arXiv:1410.4950v4 [cs.LO] 8 Nov 2015

Near-Optimal Scheduling for LTL with Future Discounting

Shota Nakagawa and Ichiro Hasuo

Department of Computer Science, The University of Tokyo

Abstract

We study the search problem for optimal schedulers for the *linear temporal logic (LTL) with future discounting*. The logic, introduced by Almagor, Boker and Kupferman, is a quantitative variant of LTL in which an event in the far future has only discounted contribution to a truth value (that is a real number in the unit interval $[0, 1]$). The precise problem we study—it naturally arises e.g. in search for a scheduler that recovers from an internal error state as soon as possible—is the following: given a Kripke frame, a formula and a number in $[0, 1]$ called a *margin*, find a path of the Kripke frame that is optimal with respect to the formula up to the prescribed margin (a truly optimal path may not exist). We present an algorithm for the problem; it works even in the extended setting with propositional quality operators, a setting where (threshold) model-checking is known to be undecidable.

1998 ACM Subject Classification F.1.1 Models of Computation

Keywords and phrases quantitative verification, optimization, temporal logic

Digital Object Identifier 10.4230/LIPIcs.xxx.yyy.p

1 Introduction

In the field of *formal methods* where a mathematical approach is taken to modeling and verifying systems, the conventional theory is built around the Boolean notion of truth: if a given system satisfies a given specification, or not. This *qualitative* theory has produced an endless list of notable achievements from hardware design to communication protocols. Among many techniques, *automata-based* ones for verification and synthesis have been particularly successful in serving engineering needs, by offering a specification method by temporal logic and push button-style algorithms. See e.g. [19, 22].

However, trends today in the use of computers—computers as part of more and more *heterogeneous* systems—have pushed researchers to turn to *quantitative* consideration of systems, too. For example, in an *embedded system* where a microcomputer controls a bigger system with mechanical/electronic components, concerns include *real-time properties*—if an expected task is finished within the prescribed deadline—and *resource consumption* e.g. with respect to electricity, memory, etc.

Quantities in formal methods can thus arise from a specification (or an *objective*) that is quantitative in nature. Another source of quantities are systems that are themselves quantitative, such as one with probabilistic behaviors.

Besides, quantities can arise simply via *refinement* of the Boolean notion of satisfaction. For example, consider the usual interpretation of the *linear temporal logic (LTL)* formula $\mathsf{F}\varphi$—it is satisfied by a sequence $s_0 s_1 \dots$ if there exists i such that $s_i \models \varphi$. It has the following natural quantitative refinement, where the modality F is replaced with a *discounted* modality $\mathsf{F}_{\exp\frac{1}{2}}$:

$$[\![s_0 s_1 \dots, \mathsf{F}_{\exp\frac{1}{2}}\varphi]\!] \;=\; (\tfrac{1}{2})^i \;, \quad \text{where } i \text{ is the least index such that } s_i \models \varphi. \tag{1}$$

© Shota Nakagawa and Ichiro Hasuo;
licensed under Creative Commons License CC-BY

Conference title on which this volume is based on.
Editors: Billy Editor and Bill Editors; pp. 1–35

Leibniz International Proceedings in Informatics
Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

This value $[\![s_0 s_1 \ldots, \mathsf{F}_{\exp \frac{1}{2}} \varphi]\!] \in [0,1]$ is a quantitative *truth value* and is like *utility* in the game-theoretic terminology. Such refinements allow quantitative reasoning about so-called *quality of service (QoS)*, specifically "how soon φ becomes true" in this example. Another example is a quantitative variation $\mathsf{G}_{\exp \frac{1}{2}} \varphi$ of $\mathsf{G}\varphi$, where $[\![s_0 s_1 \ldots, \mathsf{G}_{\exp \frac{1}{2}} \varphi]\!] = 1 - (\frac{1}{2})^i$—where i is the least index such that $s_i \not\models \varphi$—meaning that violation of φ in the far future only has a small negative impact.

$\mathrm{LTL}^{\mathrm{disc}}[\mathcal{D}, \mathcal{F}]$: LTL with Future Discounting The last examples are about quantitative refinement of temporal specifications. An important step in this direction is taken in the recent work [3]. There various useful quantitative refinements in LTL—including the last examples—are unified under the notion of *future discounting*, an idea first presented in [12] in the field of formal methods. They introduce a clean syntax of the logic $\mathrm{LTL}^{\mathrm{disc}}[\mathcal{D}, \mathcal{F}]$—called *LTL with discounting*—that combines: 1) a "discounting until" operator U_η; 2) the usual features of LTL such as the non-discounting one U; and 3) so-called propositional quality operators such as the (binary) average operator $\oplus$, in addition to $\wedge$ and $\vee$. In [3] they define its semantics; and importantly, they show that usual automata-theoretic techniques for verification and synthesis (e.g. from [19, 22]) mostly remain applicable.

Probably the most important algorithm in [3] is for the *threshold model-checking problem*: given a Kripke structure $\mathcal{K}$, a formula φ and a *threshold* $v \in [0,1]$, it asks if $[\![\mathcal{K}, \varphi]\!] > v$, i.e. the worst case truth value of a path of $\mathcal{K}$ is above v or not. The core idea of the algorithm is what we call an *event horizon*: assuming that a discounting function η in U_η tends to 0 as time goes by, and that $v > 0$, there exists a time beyond which nothing is significant enough to change the answer to the threshold model-checking problem. In this case we can approximate an infinite path by its finite prefix.

Our Contribution: Near-Optimal Scheduling for $\mathrm{LTL}^{\mathrm{disc}}[\mathcal{D}, \mathcal{F}]$ Now that a temporal formula φ assigns quantitative *truth* or *utility* $[\![\xi, \varphi]\!]$ to each path ξ, a natural task is to find a path ξ_0 in a given Kripke structure $\mathcal{K}$ that achieves the optimal. On the ground that the logic $\mathrm{LTL}^{\mathrm{disc}}[\mathcal{D}, \mathcal{F}]$ from [3] is capable of expressing many common specifications encountered in real-world problems, finding an optimal path—i.e. resolving nondeterminism in the best possible way—must have numerous applications. The situation is similar to one with *timed automata*, for which optimal scheduling problems are studied e.g. in [1].

It turns out, however, that a (truly) optimal path need not exist (Example 4.1): $v_0 = \sup_{\xi \in \mathrm{path}(\mathcal{K})} [\![\xi, \varphi]\!]$ is obviously a limit point but no ξ_0 achieves $[\![\xi_0, \varphi]\!] = v_0$. This leads us to the following *near*-optimal scheduling problem:

> **Near-optimal scheduling.** Given a Kripke structure $\mathcal{K}$, an $\mathrm{LTL}^{\mathrm{disc}}[\mathcal{D}, \mathcal{F}]$ formula φ and a *margin* $\varepsilon \in (0,1)$, find a path $\xi_0 \in \mathrm{path}(\mathcal{K})$ that is *ε-optimal*, that is, $\sup_{\xi \in \mathrm{path}(\mathcal{K})} [\![\xi, \varphi]\!] - \varepsilon \leq [\![\xi_0, \varphi]\!]$.

We study automata-theoretic algorithms for this problem. In the basic setting where there are no propositional quality operators, we can find a straightforward algorithm that conducts binary search using the model-checking algorithm from [3]. Our main contribution, however, is an alternative algorithm that takes the usual workflow: it constructs, from a formula φ and a margin ε, an automaton $\mathcal{A}_{\varphi,\varepsilon}$ with which we combine a system model $\mathcal{K}$; running a nonemptiness check-like algorithm to the resulting automaton then yields an answer.

On the one hand, our (alternative) algorithm resembles the one in [3]. In particular it relies on the idea of event horizon: a margin ε in our setting plays the role of a threshold v in [3] and enables us to ignore events in the far future.

On the other hand, a major difference from [3] is that we translate a specification (φ, ε) into an automaton that is itself quantitative (what we call a $[0,1]$-*acceptance automaton*,

with Boolean branching and $[0,1]$-acceptance values). This is unlike [3] where the target automaton is totally Boolean. An advantage of $[0,1]$-acceptance automata is that they allow optimal path search much like emptiness of Büchi automata is checked (via lasso computations). Applied to our current problem, this enables us to directly find a near-optimal path for $\mathrm{LTL}^{\mathrm{disc}}[\mathcal{D},\mathcal{F}]$ without knowing the optimal value $\sup_{\xi\in\mathrm{path}(\mathcal{K})}[\![\xi,\varphi]\!]$.

Presence of $\oplus$ and Other Propositional Quality Operators Notably, our (alternative) algorithm is shown to work even in the presence of any propositional quality operators that are *monotone* and *continuous* (in the sense we will define later; an example is the average operator $\oplus$). Those operators makes the logic more complex: indeed [3] shows that, in presence of the average operator $\oplus$, the model-checking problem for the logic $\mathrm{LTL}^{\mathrm{disc}}[\mathcal{D},\mathcal{F}]$ becomes undecidable. The binary-search algorithm mentioned earlier (that repeats model checking) ceases to work for this reason; our alternative algorithm works, nevertheless.

We analyze the complexity of the proposed algorithm, focusing on a certain subclass of the logic $\mathrm{LTL}^{\mathrm{disc}}[\mathcal{D},\mathcal{F}]$ (§4.3). Furthermore we present our prototype implementation and some experimental results (§5). They all seem to suggest the following: addition of propositional quality operators (like the average operator $\oplus$) does incur substantial computational costs—as is expected from the fact that $\oplus$ makes model checking undecidable; still our automata-theoretic approach is a viable approach, potentially applicable to optimization problems in the field of model-based system design.

The significance of the average operator $\oplus$ in envisaged applications is that it allows one to *superpose* multiple objectives. For example, one would want an event φ as soon as possible, but at the same time avoiding a different event ψ as long as possible. This is a trade-off situation and the formula $\mathsf{F}_\eta\varphi\oplus\mathsf{G}_{\eta'}\neg\psi$—with suitable discounting functions η,η'—represents a 50-50 trade-off. Other trade-off ratios can be represented as (monotone and continuous) proportional quality operators, too, and our algorithm accommodates them.

Related Work Quantitative temporal logics and their decision procedures have been a very active research topic [2,3,7,12,14]. We shall lay them out along a basic taxonomy. We denote by $\mathcal{K}$ (the model of) the system against which a specification formula φ is verified (or tested, synthesized, etc.).

- *Quantitative vs. Boolean system models.* Sometimes we need quantitative considerations just because the system $\mathcal{K}$ itself is quantitative. This is the case e.g. when $\mathcal{K}$ is a Markov chain, a Markov decision process, a timed or hybrid automaton, etc. In the current work $\mathcal{K}$ is a Kripke structure and is Boolean.
- *Quantitative vs. Boolean truth values.* The previous distinction is quite orthogonal to whether a formula φ has truth values from $[0,1]$ (or another continuous domain), or from $\{\mathtt{tt},\mathtt{ff}\}$. For example, the temporal logic PCTL [15] for reasoning about probabilistic systems has modalities like $\mathcal{P}_{>v}\psi$ ("ψ with a probability $> v$") and has Boolean interpretation. In $\mathrm{LTL}^{\mathrm{disc}}[\mathcal{D},\mathcal{F}]$ studied here, truth values are from $[0,1]$.
- *Linear time vs. branching time.* This distinction is already there in the qualitative/-Boolean setting [21]—its probabilistic variant is studied in [11]—and gives rise to temporal logics with the corresponding flavors (LTL vs. CTL, CTL^*). In fact the idea of future discounting is first introduced to a branching-time logic in [12], where an approximation algorithm for truth values is presented.
- *Future discounting vs. future averaging.* The temporal quantitative operators in $\mathrm{LTL}^{\mathrm{disc}}[\mathcal{D},\mathcal{F}]$ are *discounting*—an event's significance tends to 0 as time proceeds—a fact that benefits model checking via event horizons. Different temporal quantitative operators are studied in [7], including the *long-run average* operator $\widetilde{\mathsf{G}}\psi$. Presence of $\widetilde{G}$, however, makes most common decision problems undecidable [7].

In [14] LTL (without additional quantitative operators) is interpreted over the unit interval $[0,1]$, and its model-checking problem against quantitative systems $\mathcal{K}$ is shown to be decidable. In this setting—where the LTL connectives are interpreted by idempotent operators min and max—the variety of truth values arises only from a finite-state quantitative system $\mathcal{K}$, hence is finite.

In [3, Thm. 4] it is proved that the *threshold synthesis* problem for the logic $\mathrm{LTL}^{\mathrm{disc}}[\mathcal{D},\emptyset]$ (see Def. 2.4) is feasible. This problem asks: given a partition of atomic propositions into the input and output signals, an $\mathrm{LTL}^{\mathrm{disc}}[\mathcal{D},\emptyset]$ formula φ and $v \in [0,1]$, to come up with a transducer (i.e. a finite-state strategy) that makes the truth value of φ at least v. We remark that this is different from the near-optimal scheduling problem that we solve in this paper. The *synthesis* problem in [2, §2.2], without a threshold, is closer to ours.

Automata- (or game-) theoretic approaches are taken in [6, 8] to the synthesis of controllers or programs with better quantitative performance, too. In these papers, a specification is given itself as an automaton, instead of a temporal formula in the current work. Another difference is that, in [6,8], utility is computed along a path by limit-averaging, not future discounting. The algorithms in [6, 8] therefore rely on those which are known for mean-payoff games, including the ones in [10].

More and more diverse quantitative measures of systems' QoS are studied recently: from best/worst case probabilities and costs, to quantiles, conditional probabilities and ratios. See [5] and the references therein. Study of such in $\mathrm{LTL}^{\mathrm{disc}}[\mathcal{D},\mathcal{F}]$ is future work.

In [9] so-called *cut-point languages* of weighted automata are studied. Let $L : \Sigma^\omega \to \mathbb{R}$ be the quantitative language of a weighted automata $\mathcal{A}$. For a threshold η, the cut-point language of $\mathcal{A}$ is the set consisting of all words w such that $L(w) \geq \eta$. In [9] it is proved that the cut-point languages of deterministic limit-average automata and those of discounted-sum automata are ω-regular if the threshold η is *isolated*, that is, there is no word w such that $L(w)$ is close to η. We expect that similar properties for the logic $\mathrm{LTL}^{\mathrm{disc}}[\mathcal{D},\mathcal{F}]$ are not hard to establish, although details are yet to be worked out.

Organization of the Paper In §2 we review the logic $\mathrm{LTL}^{\mathrm{disc}}[\mathcal{D},\mathcal{F}]$ and known results on threshold model checking and satisfiability, all from [3]. We introduce quantitative variants of (alternating) Büchi automata, called (alternating) $[0,1]$-acceptance automata, in §3, with auxiliary observations on their relation to *fuzzy automata* [20]. These automata play a central role in §4 where we formalize and solve the near-optimal scheduling problem for the logic $\mathrm{LTL}^{\mathrm{disc}}[\mathcal{D},\mathcal{F}]$ (under certain assumptions on $\mathcal{D}$ and $\mathcal{F}$). We also study complexities, focusing on the average operator $\oplus$ as the only propositional quality operator. In §5 we present our implementation and some experimental results; in §6 we conclude, citing some future work. Omitted proofs are found in Appendix B.

Notations and Terminologies We shall fix some notations and terminologies, mostly following [3]. They are all standard.

The powerset of a set X is denoted by $\mathcal{P}X$. We fix the set AP of *atomic propositions*. A *computation* (over AP) is an infinite sequence $\pi = \pi_0\pi_1\ldots \in (\mathcal{P}(AP))^\omega$ over the alphabet $\mathcal{P}(AP)$. For $i \in \mathbb{N}$, $\pi^i = \pi_i\pi_{i+1}\ldots$ denotes the suffix of π starting from its i-th element.

A *Kripke structure* over AP is a tuple $\mathcal{K} = (W, R, \lambda)$ of: a finite set W of states; a transition relation $R \subseteq W^2$ that is left-total (meaning that $\forall s \in W.\, \exists s' \in W.\, (s,s') \in R$), and a labeling function $\lambda : W \to \mathcal{P}(AP)$. We follow [17] and call an infinite sequence $\xi = s_0 s_1 \ldots$ of states $s_i \in W$, such that $(s_i, s_{i+1}) \in R$ for each $i \in \mathbb{N}$, a *path* of a Kripke structure $\mathcal{K}$. The set of paths of $\mathcal{K}$ is denoted by $\mathrm{path}(\mathcal{K})$. A path $\xi = s_0 s_1 \ldots \in W^\omega$ gives rise to a computation $\lambda(s_0)\,\lambda(s_1)\ldots \in (\mathcal{P}(AP))^\omega$; the latter is denoted by $\lambda(\xi)$.

Given a set X, $\mathcal{B}^+(X)$ denotes, as usual, the set of positive propositional formulas (using

$\wedge, \vee, \top, \bot$) over $x \in X$ as atomic propositions.

2 The Logic $\mathrm{LTL}^{\mathrm{disc}}[\mathcal{D}, \mathcal{F}]$, and Its Threshold Problems

Here we recall from [2,3] our target logic, and some existing (un)decidability results.

The logic $\mathrm{LTL}^{\mathrm{disc}}[\mathcal{D}, \mathcal{F}]$ extends LTL with: 1) propositional quality operators [2] like the average operator $\oplus$; and 2) discounting in temporal operators [3]. In [3] the two extensions have been studied separately because their coexistence leads to undecidability of the (threshold) model-checking problem; here we put them altogether.

The logic $\mathrm{LTL}^{\mathrm{disc}}[\mathcal{D}, \mathcal{F}]$ has two parameters: a set $\mathcal{D}$ of discounting functions; and a set $\mathcal{F}$ of propositional connectives, called propositional quality operators.

▶ **Definition 2.1** (discounting function [3]). A *discounting function* is a strictly decreasing function $\eta : \mathbb{N} \to [0,1]$ such that $\lim_{i\to\infty} \eta(i) = 0$. A special case is an *exponential discounting function* $\exp_\lambda$, where $\lambda \in (0,1)$, that is defined by $\exp_\lambda(i) = \lambda^i$.

The set $\mathcal{D}_{\exp} = \{\exp_\lambda \mid \lambda \in (0,1) \cap \mathbb{Q}\}$ is that of exponential discounting functions.

▶ **Definition 2.2** ((monotone and continuous) propositional quality operator [2]). Let $k \in \mathbb{N}$ be a natural number. A k-ary *propositional quality operator* is a function $f : [0,1]^k \to [0,1]$.

We will eventually restrict to propositional quality operators that are *monotone* (wrt. the usual order between real numbers) and *continuous* (wrt. the usual Euclidean topology). The set of monotone and continuous propositional quality operators is denoted by $\mathcal{F}_{\mathrm{mc}}$.

▶ Example 2.3. A prototypical example of a propositional quality operator is the *average operator* $\oplus\colon [0,1]^2 \to [0,1]$, defined by $v_1 \oplus v_2 = (v_1 + v_2)/2$. (Note that $\oplus$ is a "propositional" average operator and is different from the "temporal" average operator $\widetilde{\mathsf{U}}$ in [7]). The operator $\oplus$ is monotone and continuous. Other (unary) examples from [4] include: $\triangledown_\lambda(v) = \lambda \cdot v$ and $\blacktriangledown_\lambda(v) = \lambda \cdot v + (1-\lambda)$ (they are explained in [4] to express *competence* and *necessity*, respectively). The conjunction and disjunction connectives $\wedge, \vee$, interpreted by infimums and supremums in $[0,1]$, can also be regarded as binary propositional quality operators. They are monotone and continuous, too.

Recall that the set AP is that of atomic propositions.

▶ **Definition 2.4** ($\mathrm{LTL}^{\mathrm{disc}}[\mathcal{D}, \mathcal{F}]$). Given a set $\mathcal{D}$ of discounting functions and a set $\mathcal{F}$ of propositional quality operators, the *formulas* of $\mathrm{LTL}^{\mathrm{disc}}[\mathcal{D}, \mathcal{F}]$ are defined by the grammar:

$$\varphi ::= \mathtt{True} \mid p \mid \neg\varphi \mid \varphi \wedge \varphi \mid \mathsf{X}\varphi \mid \varphi\, \mathsf{U}\, \varphi \mid \varphi\, \mathsf{U}_\eta\, \varphi \mid f(\varphi, \dots, \varphi) \ ,$$

where $p \in AP$, $\eta \in \mathcal{D}$ is a discounting function and $f \in \mathcal{F}$ is a propositional quality operator (of a suitable arity). We adopt the usual notation conventions: $\mathsf{F}\varphi = \mathtt{True}\, \mathsf{U}\, \varphi$ and $\mathsf{G}\varphi = \neg\mathsf{F}\neg\varphi$. The same goes for discounting operators: $\mathsf{F}_\eta\, \varphi = \mathtt{True}\, \mathsf{U}_\eta\, \varphi$ and $\mathsf{G}_\eta\, \varphi = \neg\, \mathsf{F}_\eta\, \neg\varphi$.

As we have already discussed, the logic $\mathrm{LTL}^{\mathrm{disc}}[\mathcal{D}, \mathcal{F}]$ extends the usual LTL with: 1) discounted temporal operators like U_η (cf. (1)); and 2) propositional quality operators like $\oplus$ that operate, on truth values from $[0,1]$ that arise from the discounted modalities, in the ways other than $\wedge$ and $\vee$ do. The precise definition below closely follows [2,3].

▶ **Definition 2.5** (semantics of $\mathrm{LTL}^{\mathrm{disc}}[\mathcal{D}, \mathcal{F}]$ [2, 3]). Let $\pi = \pi_0\pi_1 \ldots \in (\mathcal{P}(AP))^\omega$ be a computation (see §1), and φ be an $\mathrm{LTL}^{\mathrm{disc}}[\mathcal{D}, \mathcal{F}]$ formula. The *truth value* $[\![\pi, \varphi]\!]$ of φ in π

is a real number in $[0,1]$ defined as follows. Recall that $\pi^i = \pi_i\pi_{i+1}\ldots$ is a suffix of π.

$$\begin{array}{ll}
[\![\pi, \mathtt{True}]\!] = 1 & [\![\pi, p]\!] = 1 \quad (\text{if } p \in \pi_0); \qquad 0 \quad (\text{if } p \notin \pi_0) \\
[\![\pi, \neg\varphi]\!] = 1 - [\![\pi, \varphi]\!] & [\![\pi, \varphi_1 \wedge \varphi_2]\!] = \min\big\{ [\![\pi, \varphi_1]\!], [\![\pi, \varphi_2]\!] \big\} \\
[\![\pi, \mathsf{X}\varphi]\!] = [\![\pi^1, \varphi]\!] & \\
[\![\pi, \varphi_1 \,\mathsf{U}\, \varphi_2]\!] = \sup_{i\in\mathbb{N}}\big\{ \min\big\{ [\![\pi^i, \varphi_2]\!], \min_{0\le j<i} [\![\pi^j, \varphi_1]\!] \big\} \big\} & \\
[\![\pi, \varphi_1 \,\mathsf{U}_\eta\, \varphi_2]\!] = \sup_{i\in\mathbb{N}}\big\{ \min\big\{ \eta(i)[\![\pi^i, \varphi_2]\!], \min_{0\le j<i} \eta(j)[\![\pi^j, \varphi_1]\!] \big\} \big\} & \\
[\![\pi, f(\varphi_1, \ldots, \varphi_k)]\!] = f\big([\![\pi, \varphi_1]\!], \ldots, [\![\pi, \varphi_k]\!]\big) &
\end{array}$$

Compare the semantics of $\varphi_1 \,\mathsf{U}\, \varphi_2$ and that of $\varphi_1 \,\mathsf{U}_\eta\, \varphi_2$. The former is a straightforward quantitative analogue of the usual Boolean semantics; the latter additionally includes "discounting" by $\eta(i), \eta(j) \in [0,1]$. Recall that a discounting function η is deemed to be strictly decreasing; this allows us to express intuitions like in (1).

▶ **Proposition 2.6.** *The truth value* $[\![\pi, \varphi_1 \,\mathsf{U}_\eta\, \varphi_2]\!]$ *lies between* 0 *and* $\eta(0)$. ◀

We extend the semantics to Kripke structures (see §1).

▶ **Definition 2.7.** Let $\mathcal{K}$ be a Kripke structure and ξ be a path of $\mathcal{K}$. The truth value $[\![\xi, \varphi]\!]$ of φ in the path ξ is defined by $[\![\xi, \varphi]\!] = [\![\lambda(\xi), \varphi]\!]$, where $\lambda(\xi) \in (\mathcal{P}(AP))^\omega$ is the computation induced by ξ (see §1). The truth value $[\![\mathcal{K}, \varphi]\!]$ of φ in $\mathcal{K}$ is defined by $[\![\mathcal{K}, \varphi]\!] = \inf_{\xi\in\mathrm{path}(\mathcal{K})} [\![\xi, \varphi]\!]$.

▶ Remark 2.8. Later in this paper we will restrict to propositional quality operators that are monotone and continuous, i.e. $\mathrm{LTL}^{\mathrm{disc}}[\mathcal{D}, \mathcal{F}]$ with $\mathcal{F} \subseteq \mathcal{F}_{\mathrm{mc}}$. Such a logic can nevertheless express some non-monotonic operators with the help of negation. For example, the function $f_0 \colon [0,1] \to [0,1], f_0(v) = |v - \frac{1}{2}|$ can be expressed as a combination $f_0(v) = \max\{1 - f_1(v), f_2(v)\}$, using $f_1(v) = \min\{v + \frac{1}{2}, 1\}$ and $f_2(v) = \max\{v - \frac{1}{2}, 0\}$ (note that $f_1, f_2 \in \mathcal{F}_{\mathrm{mc}}$)—i.e. as the semantics of the formula $(\neg f_1\varphi) \vee (f_2\varphi)$. A nonexample is the function $f_3(v) = v \cdot \sin\frac{1}{v}$ that oscillates infinitely often in $[0,1]$.

The following "threshold" problems are studied in [3, 4]. It is shown that the logic $\mathrm{LTL}^{\mathrm{disc}}[\mathcal{D}, \emptyset]$—i.e. without propositional quality operators other than $\wedge, \vee$—has those problems decidable. Adding the average operator $\oplus$ makes them undecidable [3], while adding ∇_λ (Example 2.3) maintains decidability [4]. Here the complexities are in terms of a suitable notion $|\langle\varphi\rangle|$ of the size of φ (see [3]).

▶ **Theorem 2.9** ([3]). *The* threshold model-checking problem *for* $LTL^{\mathrm{disc}}[\mathcal{D}, \emptyset]$ *is: given a Kripke structure* $\mathcal{K}$*, an* $LTL^{\mathrm{disc}}[\mathcal{D}, \emptyset]$ *formula* φ *and a threshold* $v \in [0,1]$*, decide whether* $[\![\mathcal{K}, \varphi]\!] \ge v$*. It is decidable; when restricted to* $LTL^{\mathrm{disc}}[\mathcal{D}_{\mathrm{exp}}, \emptyset]$ *and* $v \in \mathbb{Q}$*, the problem is in PSPACE in* $|\langle\varphi\rangle|$ *and in the description of* v*, and in NLOGSPACE in the size of* $\mathcal{K}$*.*

The threshold satisfiability problem *for* $LTL^{\mathrm{disc}}[\mathcal{D}, \emptyset]$ *is: given an* $LTL^{\mathrm{disc}}[\mathcal{D}, \emptyset]$ *formula* φ*, a threshold* $v \in [0,1]$ *and* $\sim \in \{<, >\}$*, decide whether there exists a computation* $\pi \in (\mathcal{P}(AP))^\omega$ *such that* $[\![\pi, \varphi]\!] \sim v$*. This is decidable; when restricted to* $LTL^{\mathrm{disc}}[\mathcal{D}_{\mathrm{exp}}, \emptyset]$ *and* $v \in \mathbb{Q}$*, the problem is in PSPACE in* $|\langle\varphi\rangle|$ *and in the description of* v*.* ◀

▶ **Theorem 2.10** ([3]). *For* $LTL^{\mathrm{disc}}[\mathcal{D}, \{\oplus\}]$ *where* $\mathcal{D} \neq \emptyset$*, both the threshold model-checking problem and the threshold satisfiability problem are undecidable.* ◀

3 $[0,1]$-Acceptance Büchi Automata

Our algorithm for near-optimal scheduling relies on a certain notion of quantitative automaton—called $[0,1]$-*acceptance Büchi automaton*, see Def. 3.1—and an algorithm for its optimal value

problem (Lem. 3.2). The notion is not extensively studied in the literature, to the best of our knowledge.

In a $[0,1]$-acceptance Büchi automaton each state has a real value $v \in [0,1]$, instead of a Boolean value $b \in \{\mathtt{tt}, \mathtt{ff}\}$, of acceptance. Note that branching is Boolean (i.e. non-deterministic) and not $[0,1]$-weighted. In Appendix C we study a relationship to so-called *fuzzy automata* (see e.g. [20]) and show that adding weights to branching does not increase expressivity when it comes to (weighted) languages.

▶ **Definition 3.1** ($[0,1]$-acceptance automaton)**.** A $[0,1]$-*acceptance Büchi automaton*—or simply a $[0,1]$-*acceptance automaton* henceforth—is $\mathcal{A} = (\Sigma, Q, I, \delta, F)$, where Σ is a finite alphabet, Q is a finite set of states, $I \subseteq Q$ is a set of initial states, $\delta : Q \times \Sigma \to (\mathcal{P}(Q) \setminus \{\emptyset\})$ is a transition function and $F : Q \to [0,1]$ is a function that assigns an *acceptance value* to each state. We define the (weighted) language $\mathcal{L}(\mathcal{A}) : \Sigma^\omega \to [0,1]$ of $\mathcal{A}$ by

$$\mathcal{L}(\mathcal{A})(w) \;=\; \max\{F(q) \mid \exists \rho \in \mathrm{run}(w).\, q \in \mathrm{Inf}(\rho)\} \quad \text{for each } w \in \Sigma^\omega \;, \tag{2}$$

where the sets $\mathrm{run}(w)$ and $\mathrm{Inf}(\rho)$ are defined as usual. Precisely:

- For an infinite word $w \in \Sigma^\omega$, a *run* over w of $\mathcal{A}$ is an infinite alternating sequence $\rho = q_0 a_0 q_1 a_1 \dots$ such that: 1) $q_i \in Q$ is a state and $a_i \in \Sigma$ is a letter, for all $i \in \mathbb{N}$; 2) $q_0 \in I$; and 3) $q_{i+1} \in \delta(q_i, a_i)$ for all $i \in \mathbb{N}$. The set of runs over w is denoted by $\mathrm{run}(w)$.
- Given a run ρ, the set $\mathrm{Inf}(\rho)$ is defined by $\mathrm{Inf}(\rho) = \{q \in Q \mid q \text{ occurs infinitely often in } \rho\}$.

Note that, when we restrict to Boolean acceptance values (i.e. $F(q) \in \{0,1\}$), the acceptance value in (2) precisely coincides with the one in the usual notion of Büchi automaton. Note also that, in (2), we take the maximum of finitely many values (the state space Q is finite).

The following observation, though not hard, is a key fact for our search algorithm. It is a quantitative analogue of emptiness check in usual (Boolean) automata.

▶ **Lemma 3.2** (the optimal value problem for $[0,1]$-acceptance automata)**.** *Let $\mathcal{A} = (\Sigma, Q, I, \delta, F)$ be a $[0,1]$-acceptance Büchi automaton. There exists the maximum $\max_{w \in \Sigma^\omega} \mathcal{L}(\mathcal{A})(w)$ of $\mathcal{L}(\mathcal{A})$. Moreover, there is an algorithm that computes the value $\max_{w \in \Sigma^\omega} \mathcal{L}(\mathcal{A})(w)$ as well as a run $\rho_{\max} = q_0 a_0 q_1 a_1 \dots \in (\Sigma \times Q)^\omega$ that realizes the maximum.*

Proof. The algorithm is much like the one for emptiness check of (ordinary) Büchi automata, searching for a suitable lasso computation. More concretely: consider those states q which are both reachable from some initial state and reachable from q itself. Let s be one, among those states, with the greatest acceptance value $F(s)$. It is easy to show that a lasso computation with the state s as a "knot" gives the run $\rho_{\max}$ that we seek for. ◀

Our algorithm first translates a formula into an *alternating* $[0,1]$-acceptance automata.

▶ **Definition 3.3** (alternating $[0,1]$-acceptance automaton)**.** An *alternating* $[0,1]$-*acceptance (Büchi) automaton* is a tuple $\mathcal{A} = (\Sigma, Q, I, \delta, F)$, where Σ is a finite alphabet, Q is a finite set of states, $I \subseteq Q$ is a set of initial states, $\delta : Q \times \Sigma \to \mathcal{B}^+(Q \cup [0,1])$ is a transition function and $F : Q \to [0,1]$ gives acceptance values. Recall (§1) that $\mathcal{B}^+(Q \cup [0,1])$ is the set of positive propositional combinations of $q \in Q$ and $v \in [0,1]$.

We define the (weighted) language $\mathcal{L}(\mathcal{A}) : \Sigma^\omega \to [0,1]$ of $\mathcal{A}$ by

$$\mathcal{L}(\mathcal{A})(w) \;=\; \max_{\tau \in \mathrm{run}_{\mathcal{A}}(w)} \min_{\rho \in \mathrm{path}_{\mathcal{A},w}(\tau)} F^\infty(\rho) \;, \tag{3}$$

where *runs*, *paths* and the function F^∞ are formally defined much like with the usual alternating automata. Precisely:

- A run is much like with the usual alternating automata. Precisely, let $\mathcal{A} = (\Sigma, Q, I, \delta, F)$ be an alternating $[0,1]$-acceptance automaton and $w = a_0 a_1 \ldots \in \Sigma^\omega$ be an infinite word. A *run* τ of $\mathcal{A}$ over w is a (possibly infinite-depth) tree subject to the following.
  - Each node t of the tree τ is labeled from $Q \cup [0,1]$.
  - The root of τ is labeled with an initial state $q_0 \in I$.
  - Any node t labeled with a number $v \in [0,1]$ is a leaf.
  - Consider an arbitrary node t that is labeled with a state $q \in Q$. Assume that t is of depth $i \in \mathbb{N}$; and let the labels of t's children be $l_1, \ldots, l_k \in Q \cup [0,1]$. We require $l_1, \ldots, l_k \models \delta(q, a_i)$, where: $\delta(q, a_i) \in \mathcal{B}^+(Q \cup [0,1])$ is the a_i-successor of q in $\mathcal{A}$; and $\models$ designates the obvious Boolean notion of satisfaction (where we think of elements of $Q \cup [0,1]$ as atomic variables).

  The set $\mathrm{run}_{\mathcal{A}}(w)$ is that of all runs of $\mathcal{A}$ over the word w.
- A *path* ρ of a run $\tau \in \mathrm{run}_{\mathcal{A}}(w)$ is simply a (finite or infinite) path in the tree τ, from the root of τ. A path ρ is finite only when its last state is a leaf of τ. The set of paths of $\tau \in \mathrm{run}_{\mathcal{A}}(w)$ is denoted by $\mathrm{path}_{\mathcal{A},w}(\tau)$.
- The function $F^\infty \colon \mathrm{path}_{\mathcal{A},w}(\tau) \to [0,1]$ in (3) is defined as follows. If $\rho \in \mathrm{path}_{\mathcal{A},w}(\tau)$ is an infinite path, each node t in ρ is labeled with a state q of $\mathcal{A}$. We define

$$F^\infty(\rho) \;=\; \max\{\, F(q) \mid q \in Q \text{ occurs infinitely often, as labels, in } \rho \,\} \;. \tag{4}$$

  Assume now that $\rho \in \mathrm{path}_{\mathcal{A},w}(\tau)$ is finite, say $\rho = t_0 t_1 \ldots t_i$. Then the last node t_i is labeled either by $v \in [0,1]$ or $q \in [0,1]$. In the former case we define $F^\infty(\rho) = v$ (i.e. F^∞ returns the label of t_i). In the latter case, we have that $\delta(t_i, a_i)$ is propositionally equivalent to $\top$ ("truth") by the definition of run. We define $F^\infty(\rho) = 1$.

In the above we used max and min (not sup or inf) since $\{F(q) \mid q \in Q\}$ is a finite set.

▶ **Proposition 3.4.** *Let $\mathcal{A} = (\Sigma, Q, I, \delta, F)$ be an alternating $[0,1]$-acceptance automaton. There exists a $[0,1]$-acceptance automaton $\mathcal{A}'$ such that $\mathcal{L}(\mathcal{A}) = \mathcal{L}(\mathcal{A}')$.* ◀

The construction of $\mathcal{A}'$ is a quantitative adaptation of the one in [18] that turns an alternating ω-automaton into a nondeterministic one. In our adaptation we use what we call *exposition flags*, an idea that is potentially useful in other settings with Büchi-type acceptance conditions, too. See Appendix B.1 for details of the proof and the construction therein.

Later we will also use the fact that $[0,1]$-acceptance automata are closed under monotone propositional quality operators (Def. 2.2).

▶ **Proposition 3.5.** *Let $f \colon [0,1]^k \to [0,1]$ be monotone, and $\mathcal{A}_1, \ldots, \mathcal{A}_k$ be $[0,1]$-acceptance automata over a common alphabet Σ. There is a $[0,1]$-acceptance automaton $f(\mathcal{A}_1, \ldots, \mathcal{A}_k)$ such that $\mathcal{L}\big(f(\mathcal{A}_1, \ldots, \mathcal{A}_k)\big)(w) = f\big(\mathcal{L}(\mathcal{A}_1)(w), \ldots, \mathcal{L}(\mathcal{A}_k)(w)\big)$ for each $w \in \Sigma^\omega$.* ◀

▶ Remark 3.6. Prop. 3.4 and 3.5 are essentially two separate constructions that deal with: the connectives $\wedge$ and $\vee$; and the other propositional quality operators, respectively. One can alternatively think of $\wedge$ and $\vee$ as special cases of the latter (Example 2.3) and use Prop. 3.5 altogether. This however results in a worse complexity: the powerset-like construction in Prop. 3.4 exploits the commutativity, idempotency and associativity of $\wedge$ to suppress the number of states, while such is not done in the product-like construction in Prop. 3.5.

A generalization of $[0,1]$-acceptance automaton is naturally obtained by making transitions also $[0,1]$-weighted. The result is called *fuzzy automaton* and studied e.g. in [20]. In Appendix C we show that this generalization does not add expressivity. In fact we prove a more general result there, parametrizing $[0,1]$ into a suitable semiring $\mathbb{K}$.

4 Near-Optimal Scheduling for $\mathrm{LTL}^{\mathrm{disc}}[\mathcal{D}, \mathcal{F}_{\mathrm{mc}}]$

In [3,4] the threshold model-checking problem for the logic $\mathrm{LTL}^{\mathrm{disc}}[\mathcal{D}, \mathcal{F}]$ is studied. In this paper, instead, we are interested in the following problem: what path of a given Kripke structure $\mathcal{K}$ is the best for a given $\mathrm{LTL}^{\mathrm{disc}}[\mathcal{D}, \mathcal{F}]$ formula φ.

In general, however, there does not exist an optimal path ξ_0 of $\mathcal{K}$, i.e. one that achieves $[\![\xi_0, \varphi]\!] = \sup_{\xi \in \mathrm{path}(\mathcal{K})} [\![\xi, \varphi]\!]$.

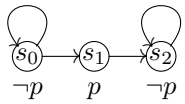

▶ Example 4.1 (optimality not achievable). Take a formula $\varphi = \mathsf{G}_\eta \mathsf{F} p$ and the Kripke structure $\mathcal{K}$ shown in the above. This example illustrates that the existence of an optimal path is not guaranteed in general: indeed, whereas $\sup_{\xi' \in \mathrm{path}(\mathcal{K})} [\![\xi', \varphi]\!] = 1$ in this example, there is no path ξ that achieves $[\![\xi, \varphi]\!] = 1$.

More specifically: we first note that, in each path ξ of the Kripke structure, p is true at most once. The later the state s_1 occurs in a path ξ, the bigger the truth value $[\![\xi, \varphi]\!]$ is; moreover the value $[\![\xi, \varphi]\!]$ tends to 1 (since η tends to 0). However there is no path ξ that achieves exactly $[\![\xi, \varphi]\!] = 1$: if p is postponed indefinitely, no state in ξ satisfies p, in which case $\mathsf{F} p$ is everywhere false and hence $[\![\xi, \varphi]\!] = 0$.

We thus strive for *near*-optimality, allowing a prescribed margin ε.

▶ **Definition 4.2.** The *near-optimal scheduling* problem for $\mathrm{LTL}^{\mathrm{disc}}[\mathcal{D}, \mathcal{F}]$ is: given a Kripke structure $\mathcal{K} = (W, R, \lambda)$, an $\mathrm{LTL}^{\mathrm{disc}}[\mathcal{D}, \mathcal{F}]$ formula φ and a positive real number $\varepsilon \in (0, 1)$, to find a path $\xi_0 \in \mathrm{path}(\mathcal{K})$ such that $[\![\xi_0, \varphi]\!] \geq \sup_{\xi \in \mathrm{path}(\mathcal{K})} [\![\xi, \varphi]\!] - \varepsilon$.

Ultimately we will show that the problem in the above is decidable (Thm. 4.14), when all the propositional quality operators are monotone and continuous ($\mathcal{F} \subseteq \mathcal{F}_{\mathrm{mc}}$).

We first note that, in the special case for $\mathrm{LTL}^{\mathrm{disc}}[\mathcal{D}, \emptyset]$ (i.e. no propositional quality operators), there is a straightforward binary search algorithm that relies on the (threshold) model-checking algorithm in [3] (Thm. 2.9). Specifically, the binary search algorithm repeatedly conducts threshold model-checking for: the threshold $v = \frac{1}{2}$ in the first round; $v = \frac{1}{4}$ or $\frac{3}{4}$ in the second round, depending on the outcome of the first round; then for $v = \frac{1}{8}, \dots, \frac{6}{8}$ or $\frac{7}{8}$, depending on the outcome of the second round; and so on. Given a margin $\varepsilon \in (0, 1)$, this way, we need $-\log \varepsilon$ rounds. This binary search algorithm is rather effective (see §5).

However the binary search algorithm does not work in presence of the average operator $\oplus$, simply because the threshold model-checking problem is undecidable (Thm. 2.10). Our main contribution is a novel algorithm for near-optimal scheduling that works even in this case (and more generally for the logic $\mathrm{LTL}^{\mathrm{disc}}[\mathcal{D}, \mathcal{F}_{\mathrm{mc}}]$). Our algorithm first translates a formula φ and a margin $\varepsilon \in (0, 1)$ to an alternating $[0, 1]$-acceptance automaton $\mathcal{A}_{\varphi, \varepsilon}$, which is further turned into a $[0, 1]$-acceptance automaton (Prop. 3.4). The resulting automaton—after taking the product with $\mathcal{K}$—is amenable to optimal value search (Lem. 3.2), yielding a solution to the original problem.

In the rest of the section we describe our algorithm. We shall however first restrict to the logic $\mathrm{LTL}^{\mathrm{disc}}[\mathcal{D}, \emptyset]$ for the sake of presentation (although this basic fragment allows binary search). After describing the basic algorithm for $\mathrm{LTL}^{\mathrm{disc}}[\mathcal{D}, \emptyset]$ in §4.1, in §4.2 we explain how it can be modified to accommodate propositional quality operators.

4.1 Our Algorithm, When Restricted to $\mathrm{LTL}^{\mathrm{disc}}[\mathcal{D}, \emptyset]$

Our translation of φ and $\varepsilon \in (0, 1)$ to an automaton $\mathcal{A}_{\varphi, \varepsilon}$ is an extension of the standard translation from LTL formulas to alternating Büchi automata (e.g. in [22]), with:

- incorporation of quantities—accumulation of discount factors, more specifically—by means of what we call *discount sequences*; and

- cutting off those events which are far in the future—the idea of *event horizon* from [3].

The extension is not complicated on the conceptual level. Its details need care, however, especially in handling negations and alternation of greatest and least fixed points.

As preparation, we recall some definitions and notations from [3].

▶ **Definition 4.3** (η^{+k}, $xcl(\varphi)$ [3]). Let $\eta : \mathbb{N} \to [0,1]$ be a discounting function. We define a discounting function $\eta^{+k} : \mathbb{N} \to [0,1]$ by $\eta^{+k}(i) = \eta(i+k)$ for each $k \in \mathbb{N}$.

For an LTL$^{\text{disc}}[\mathcal{D}, \mathcal{F}]$ formula φ, the *extended closure* $xcl(\varphi)$ of φ [3] is defined by

$$xcl(\varphi) \;=\; \mathrm{Sub}(\varphi) \cup \{\varphi_1 \,\mathsf{U}_{\eta^{+k}}\, \varphi_2 \mid k \in \mathbb{N}, \varphi_1 \,\mathsf{U}_{\eta}\, \varphi_2 \in \mathrm{Sub}(\varphi)\} \;, \tag{5}$$

where $\mathrm{Sub}(\varphi)$ denotes the set of subformulas of φ.

4.1.1 Discounting Sequences

We go on to technical details. In the alternating $[0,1]$-acceptance automaton $\mathcal{A}_{\varphi,\varepsilon}$ that we shall construct, a state is a pair $(\psi, \vec{d})$ of a formula ψ and a *discount sequence* $\vec{d} \in [0,1]^+$.

▶ **Definition 4.4** (discount sequence). A *discount sequence* is a sequence $\vec{d} = d_1 d_2 \dots d_n \in [0,1]^+$ of real numbers with a nonzero length ($d_i \in [0,1]$ for each i).

The notion of discount sequence is a quantitative extension of that of *priority* in parity automata. Specifically, the length n of a discount sequence $\vec{d} = d_1 d_2 \dots d_n$ corresponds to a priority—i.e. the alternation depth of greatest and least fixed points. Each real number d_i in the sequence, in turn, stands for the accumulated discount factor in each level of fixed-point alternation. For example, the formula $\mathsf{F}_{\exp \frac{1}{2}}\, \mathsf{G}_{\exp \frac{2}{3}}\, \mathsf{F}_{\exp \frac{3}{4}}\, p$ will induce a discount sequence $(\frac{1}{2})^{n_1}, (\frac{2}{3})^{n_2}, (\frac{3}{4})^{n_3}$ of length 3—where n_1, n_2 and n_3 are the numbers of steps for which the three discounting temporal operators F_{η_1}, G_{η_2} and F_{η_3} "have waited," respectively.

We use three operators $\odot, :, \boxtimes$ that act on discount sequences; the intuitions are as follows. The first two are for accumulating discount factors: we use $\odot$ in case there is no alternation of greatest and least fixed points; and we use : in case there is. Examples are:

$$\begin{array}{lcl} \big((\frac{1}{2})^2, (\frac{2}{3})^3, \frac{3}{4}\big) \odot \frac{4}{5} & = & \big((\frac{1}{2})^2, (\frac{2}{3})^3, \frac{3}{4} \cdot \frac{4}{5}\big) \;=\; \big((\frac{1}{2})^2, (\frac{2}{3})^3, \frac{3}{5}\big) \quad \text{and} \\ \big((\frac{1}{2})^2, (\frac{2}{3})^3, \frac{3}{4}\big) : \frac{4}{5} & = & \big((\frac{1}{2})^2, (\frac{2}{3})^3, \frac{3}{4}, \frac{4}{5}\big) \;. \end{array}$$

Note that in the former the length is preserved, while in the latter the sequence gets longer by one.

▶ **Definition 4.5** ($\vec{d} \odot d'$, $\vec{d} : d'$). The operator $\odot$ takes a discount sequence $\vec{d}$ and a discount factor $d' \in [0,1]$ as arguments, and multiplies the last element of $\vec{d}$ by d'. That is,

$$(d_1 d_2 \dots d_n) \odot d' \;=\; d_1 d_2 \dots d_{n-1}(d_n \cdot d') \quad \in [0,1]^+ \;.$$

The operator : is simply the concatenation operator: given $\vec{d} = d_1 d_2 \dots d_n$ and $d' \in [0,1]$, the sequence $\vec{d} : d'$ is $d_1 d_2 \dots d_n d'$ of length $n+1$.

We use the operator $\boxtimes$ in $\vec{d} \boxtimes v$ to let a discount sequence $\vec{d}$ act on a truth value $v \in [0,1]$.

▶ **Definition 4.6** ($\vec{d} \boxtimes v$). The operator $\boxtimes$ takes $\vec{d} \in [0,1]^+$ and $v \in [0,1]$ as arguments. The value $\vec{d} \boxtimes v \in [0,1]$ is defined inductively by:

$$d \boxtimes v \;=\; dv \;, \qquad \vec{d}d' \boxtimes v \;=\; \vec{d} \boxtimes (1 - d'v) \;. \quad \text{Explicitly:} \tag{6}$$

$$(d_1 d_2 \dots d_n) \boxtimes v \;=\; d_1 \,-\, d_1 d_2 \,+\, d_1 d_2 d_3 \,-\, \cdots \,+\, (-1)^n d_1 d_2 \dots d_{n-1} \,+\, (-1)^{n+1} d_1 d_2 \dots d_n v \,. \tag{7}$$

The intuition behind the action $\vec{d} \boxtimes v$ is most visible in (6), where dv and $d'v$ denote multiplication of real numbers. Given a discount sequence $\vec{d}d'$: 1) we apply the final discount factor d' to the truth value v, obtaining $d'v$; 2) the alternation between greatest and least fixed points is taken into account, by taking the negation $1 - d'v$ (cf. Def. 2.5); and 3) we apply the remaining sequence $\vec{d}$ inductively and obtain $\vec{d} \boxtimes (1 - d'v)$. An example is $\left(\frac{3}{4}, \frac{1}{3}, \frac{2}{5}\right) \boxtimes 1 = \left(\frac{3}{4}, \frac{1}{3}\right) \boxtimes \left(1 - \frac{2}{5} \cdot 1\right) = \left(\frac{3}{4}, \frac{1}{3}\right) \boxtimes \frac{3}{5} = \left(\frac{3}{4}\right) \boxtimes \left(1 - \frac{1}{3} \cdot \frac{3}{5}\right) = \left(\frac{3}{4}\right) \boxtimes \frac{4}{5} = \frac{3}{5}$.

The following relationship between $\odot$ and $\boxtimes$ is easily seen to hold:

$$(\vec{d} \odot d') \boxtimes v \;=\; \vec{d} \boxtimes (d' \cdot v) \;. \tag{8}$$

The three operators $\odot, :, \boxtimes$ defined in the above will be used shortly, in the construction of the automaton $\mathcal{A}_{\varphi,\varepsilon}$. Their roles are briefly discussed after Def. 4.7.

4.1.2 Construction of $\mathcal{A}_{\varphi,\varepsilon}$

We describe the construction of $\mathcal{A}_{\varphi,\varepsilon}$, for a formula φ of $\mathrm{LTL}^{\mathrm{disc}}[\mathcal{D}, \emptyset]$ and a margin ε. We subsequently discuss ideas behind it, comparing the definition with other known constructions.

We first define $\mathcal{A}^{\mathrm{p}}_{\varphi,\varepsilon}$ that is infinite-state, and obtain $\mathcal{A}_{\varphi,\varepsilon}$ as the reachable part. The latter will be shown to be finite-state (Lem. 4.8).

▶ **Definition 4.7** (the automata $\mathcal{A}^{\mathrm{p}}_{\varphi,\varepsilon}, \mathcal{A}_{\varphi,\varepsilon}$)**.** Let φ be an $\mathrm{LTL}^{\mathrm{disc}}[\mathcal{D}, \emptyset]$ formula and $\varepsilon \in (0, 1)$. We define an alternating $[0, 1]$-acceptance automaton $\mathcal{A}^{\mathrm{p}}_{\varphi,\varepsilon} = (\mathcal{P}(AP), Q, I, \delta, F)$ as follows. Its state space Q is $xcl(\varphi) \times [0, 1]^+$; hence a state is a pair $(\psi, \vec{d})$ of a formula and a discount sequence. The transition function $\delta \colon Q \times \mathcal{P}(AP) \to \mathcal{B}^+(Q \cup [0, 1])$ is defined as in Table 1, where we let $\vec{d} = d_1 d_2 \dots d_n \in [0, 1]^+$ and $\sigma \in \mathcal{P}(AP)$.

The set I of the initial states of $\mathcal{A}^{\mathrm{p}}_{\varphi,\varepsilon}$ is $\{(\varphi, 1)\}$. The acceptance function F is

$$F(\psi, \vec{d}) = \begin{cases} 1 & \text{if } \psi = \psi_1 \,\mathsf{U}\, \psi_2 \text{ and } |\vec{d}| \text{ is even} \\ 0 & \text{otherwise.} \end{cases} \tag{12}$$

The alternating $[0, 1]$-acceptance automaton $\mathcal{A}_{\varphi,\varepsilon}$ is defined to be the restriction of $\mathcal{A}^{\mathrm{p}}_{\varphi,\varepsilon}$ to the states that are reachable from the initial state $(\varphi, 1)$.

Examples of $\mathcal{A}_{\varphi,\varepsilon}$ are in Fig. 1–2, where $(\varphi, \varepsilon) = (\mathsf{G}_{\exp \frac{1}{2}}\, \mathsf{F}_{\exp \frac{2}{5}}\, p, \frac{1}{3})$ and $(\mathsf{F}_{\exp \frac{1}{2}}\, \mathsf{G} p, \frac{1}{3})$. There a discount sequence $d_1 \dots d_n$ is denoted by $\langle d_1, \dots, d_n \rangle$ for readability.

Some remarks on Def. 4.7 are in order.

In Absence of Discounting (Sanity Check) If the formula φ contains no discounting operator U_η, then the construction essentially coincides the usual one in [22] that translates a (usual) LTL formula to an alternating Büchi automaton. To see it, recall that the length $|\vec{d}|$ of a discount sequence plays the role of a priority in parity automata (§4.1.1). Therefore in the first case of (12), $|\vec{d}|$ being even means that we are in fact dealing with a greatest fixed point. This makes the state accepting (in the Büchi sense), much like in [22].

$\mathcal{A}_{\varphi,\varepsilon}$ is Quantitative The acceptance values of the states of $\mathcal{A}_{\varphi,\varepsilon}$ are Boolean (see (12)). Nevertheless the automaton is quantitative, in that non-Boolean values from $[0, 1]$ appear as atomic propositions in the range $\mathcal{B}^+(Q \cup [0, 1])$ of the transition δ (they occur at the leaves in Fig. 1–2). Once we transform $\mathcal{A}_{\varphi,\varepsilon}$ to a non-alternating automaton (Prop. 3.4), these non-Boolean propositional values give rise to non-Boolean acceptance values.

Event Horizon A fundamental idea from [3] is the following. A discounting operator, in presence of a threshold (in [3]) or a nonzero margin (here), allows an exact representation

$$
\begin{aligned}
\delta\big((\mathtt{True}, \vec{d}), \sigma\big) &= \vec{d} \boxtimes 1 \\
\delta\big((p, \vec{d}), \sigma\big) &= \begin{cases} \vec{d} \boxtimes 1 & \text{if } p \in \sigma, \\ \vec{d} \boxtimes 0 & \text{otherwise.} \end{cases} \\
\delta\big((\neg\psi, \vec{d}), \sigma\big) &= \delta\big((\psi, \vec{d}\colon 1), \sigma\big) \\
\delta\big((\psi_1 \wedge \psi_2, \vec{d}), \sigma\big) &= \begin{cases} \delta\big((\psi_1, \vec{d}), \sigma\big) \wedge \delta\big((\psi_2, \vec{d}), \sigma\big) & \text{if } |\vec{d}| \text{ is odd,} \\ \delta\big((\psi_1, \vec{d}), \sigma\big) \vee \delta\big((\psi_2, \vec{d}), \sigma\big) & \text{otherwise.} \end{cases} \\
\delta\big((\mathsf{X}\psi, \vec{d}), \sigma\big) &= (\psi, \vec{d}) \\
\delta\big((\psi_1 \,\mathsf{U}\, \psi_2, \vec{d}), \sigma\big) &= \begin{cases} \delta\big((\psi_2, \vec{d}), \sigma\big) \vee \Big(\delta\big((\psi_1, \vec{d}), \sigma\big) \wedge (\psi_1 \,\mathsf{U}\, \psi_2, \vec{d})\Big) & \text{if } |\vec{d}| \text{ is odd,} \\ \delta\big((\psi_2, \vec{d}), \sigma\big) \wedge \Big(\delta\big((\psi_1, \vec{d}), \sigma\big) \vee (\psi_1 \,\mathsf{U}\, \psi_2, \vec{d})\Big) & \text{otherwise.} \end{cases}
\end{aligned} \tag{9}
$$

For $\delta\big((\psi_1 \,\mathsf{U}_\eta\, \psi_2, \vec{d}), \sigma\big)$ we make cases. Let $\vec{d} = d_1 \dots d_n$. If $\eta(0) \cdot \prod_{i=1}^{n} d_i \leq \varepsilon$:

$$
\delta\big((\psi_1 \,\mathsf{U}_\eta\, \psi_2, \vec{d}), \sigma\big) = \begin{cases} \vec{d} \boxtimes 0 & \text{if } |\vec{d}| \text{ is odd,} \\ \vec{d} \boxtimes \eta(0) & \text{otherwise;} \end{cases} \tag{10}
$$

otherwise, i.e. if $\eta(0) \cdot \prod_{i=1}^{n} d_i > \varepsilon$:

$$
\delta\big((\psi_1 \,\mathsf{U}_\eta\, \psi_2, \vec{d}), \sigma\big) = \begin{cases} \delta\big((\psi_2, \vec{d} \odot \eta(0)), \sigma\big) \vee \Big(\delta\big((\psi_1, \vec{d} \odot \eta(0)), \sigma\big) \wedge (\psi_1 \,\mathsf{U}_{\eta+1}\, \psi_2, \vec{d})\Big) & \text{if } |\vec{d}| \text{ is odd,} \\ \delta\big((\psi_2, \vec{d} \odot \eta(0)), \sigma\big) \wedge \Big(\delta\big((\psi_1, \vec{d} \odot \eta(0)), \sigma\big) \vee (\psi_1 \,\mathsf{U}_{\eta+1}\, \psi_2, \vec{d})\Big) & \text{otherwise.} \end{cases} \tag{11}
$$

Table 1 Transition function δ of $\mathcal{A}^{\mathrm{p}}_{\varphi,\varepsilon}$

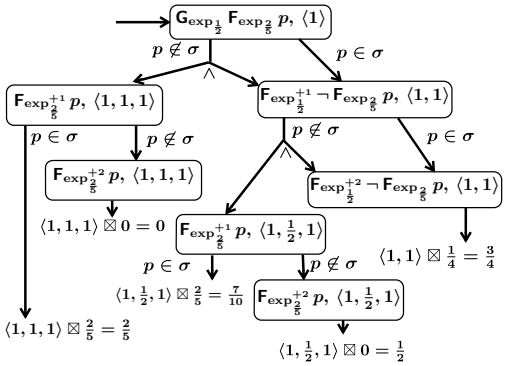

Figure 1 The automaton $\mathcal{A}_{\varphi,\varepsilon}$ for $\varphi = \mathsf{G}_{\exp\frac{1}{2}}\,\mathsf{F}_{\exp\frac{2}{5}}\,p$ and $\varepsilon = \frac{1}{3}$

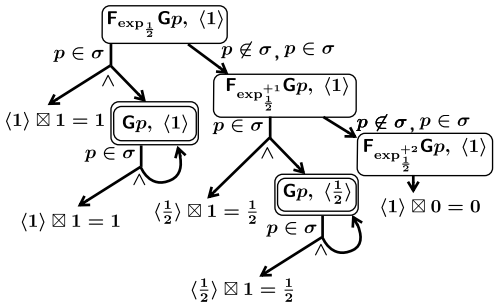
$\mathsf{F}_{\exp\frac{1}{2}}\,\mathsf{G}p,\ \langle 1\rangle$
$p \in \sigma$
$p \notin \sigma, p \in \sigma$
$\langle 1\rangle \boxtimes 1 = 1$
$\mathsf{G}p,\ \langle 1\rangle$
$\mathsf{F}_{\exp^{+1}_{\frac{1}{2}}}\,\mathsf{G}p,\ \langle 1\rangle$
$p \in \sigma$
$p \notin \sigma, p \in \sigma$
$p \in \sigma$
$\mathsf{F}_{\exp^{+2}_{\frac{1}{2}}}\,\mathsf{G}p,\ \langle 1\rangle$
$\langle 1\rangle \boxtimes 1 = 1$
$\mathsf{G}p,\ \langle \frac{1}{2}\rangle$
$\langle 1\rangle \boxtimes 0 = 0$
$\langle \frac{1}{2}\rangle \boxtimes 1 = \frac{1}{2}$
$p \in \sigma$
$\langle \frac{1}{2}\rangle \boxtimes 1 = \frac{1}{2}$

Figure 2 The automaton $\mathcal{A}_{\varphi,\varepsilon}$ for $\varphi = \mathsf{F}_{\exp\frac{1}{2}}\,\mathsf{G}p$ and $\varepsilon = \frac{1}{3}$. The double-lined nodes have the acceptance value 1.

by a (finitary) formula without a fixed point operator. The latter means, for example:

$$[\![\pi, \mathsf{F}_{\exp \frac{1}{2}}\, \varphi]\!] \geq \tfrac{1}{4} \iff \pi \models \varphi \vee \mathsf{X}\varphi \vee \mathsf{XX}\varphi \ , \quad \text{and} \tag{13}$$

$$[\![\pi, \mathsf{G}_{\exp \frac{1}{2}}\, \varphi]\!] \geq \tfrac{3}{4} \iff \pi \models \varphi \wedge \mathsf{X}\varphi \wedge \mathsf{XX}\varphi \ , \tag{14}$$

and so on. Note that in (13), whatever happens after two time units has contributions less than $(\frac{1}{2})^2 = \frac{1}{4}$ and therefore never enough to make up the threshold. The example (14) is similar, with events in the future having only negligible negative contributions. In other words: fixed point operators with discounting have an *event horizon*—in the above examples (13–14) it lies between $t = 2$ and 3—nothing beyond which matters.

This idea of event horizon is used in the distinction between (10) and (11). The value $\eta(0) \cdot \prod_{i=1}^{n} d_i$ is, as we shall see, the greatest contribution to a truth value that the events henceforth potentially have. In case it is smaller than the margin ε we can safely ignore the *positive* contribution henceforth and take the smallest possible truth value 0—much like the disjunct $\mathsf{X}^3\varphi \vee \mathsf{X}^4\varphi \vee \cdots$ is truncated in (13). This is what is done in the first case in (10). The second case in (10) is about a greatest fixed point and we truncate the *negative* contributions of the events beyond the event horizon—this is much like the obligation $\mathsf{X}^3\varphi \wedge \mathsf{X}^4\varphi \wedge \cdots$ is lifted in (14). In this case we use the greatest truth value possible, namely $\eta(0)$. This is what is done in (10).

Use of Discount Sequences Discount sequences $\vec{d}$ are used for two purposes. Firstly, as we already described, its length $|\vec{d}|$ indicates the alternation between positive and negative views on a formula—observe that a discount sequence gets longer in (9). Consequently many clauses in the definition of δ distinguish cases according to the parity of $|\vec{d}|$. Secondly it records all the discount factors that have been encountered. See (11), where the last element of $\vec{d}$ is multiplied by the newly encountered factor $\eta(0)$ and updated to $\vec{d} \odot \eta(0)$. Such accumulation $\vec{d}$ of discount factors acts on a truth value via the $\boxtimes$ operator, like in (10) and in the definition of $\delta\big((\mathtt{True}, \vec{d}), \sigma\big)$.

▶ **Lemma 4.8.** *The automaton $\mathcal{A}_{\varphi,\varepsilon}$ has only finitely many states.* ◀

The following "correctness lemma" claims that $\mathcal{A}_{\varphi,\varepsilon}$ conducts the expected task.

▶ **Lemma 4.9.** *Let φ be an $LTL^{\mathrm{disc}}[\mathcal{D}, \emptyset]$ formula and $\varepsilon \in (0,1)$ be a positive real number. For each computation $\pi \in (\mathcal{P}(AP))^\omega$, we have $[\![\pi, \varphi]\!] - \varepsilon \ \leq \ \mathcal{L}(\mathcal{A}_{\varphi,\varepsilon})(\pi) \ \leq \ [\![\pi, \varphi]\!]$* ◀

4.1.3 The Algorithm

After the construction of $\mathcal{A}_{\varphi,\varepsilon}$, the algorithm proceeds in the following manner. We first translate $\mathcal{A}_{\varphi,\varepsilon}$ to a (non-alternating) $[0,1]$-acceptance automaton (relying on Prop. 3.4).

▶ **Corollary 4.10.** *Let φ be an $LTL^{\mathrm{disc}}[\mathcal{D}, \emptyset]$ formula and $\varepsilon \in (0,1)$ be a positive real number. There exists a (non-alternating) $[0,1]$-acceptance automaton $\mathcal{A}^{\mathrm{na}}_{\varphi,\varepsilon}$ such that $[\![\pi, \varphi]\!] - \varepsilon \leq \mathcal{L}(\mathcal{A}^{\mathrm{na}}_{\varphi,\varepsilon})(\pi) \leq [\![\pi, \varphi]\!]$ for each computation $\pi \in (\mathcal{P}(AP))^\omega$.* ◀

Towards the solution of the near-optimal scheduling problem (Def. 4.2), we construct the *product* of $\mathcal{A}^{\mathrm{na}}_{\varphi,\varepsilon}$ in Cor. 4.10 and the given Kripke structure $\mathcal{K}$. Since transitions of $[0,1]$-acceptance automata are nondeterministic, this product can be defined just as usual.

▶ **Definition 4.11.** Let $\mathcal{A} = (\mathcal{P}(AP), Q, I, \delta, F)$ be a $[0,1]$-acceptance automaton and $\mathcal{K} = (W, R, \lambda)$ be a Kripke structure. Their *product* $\mathcal{A} \times \mathcal{K}$ is a $[0,1]$-acceptance automaton $(1, Q', I', \delta', F')$—over a singleton alphabet $1 = \{\bullet\}$—defined by: $Q' = Q$; $I' = I \times W$; $\delta'\big((q,s), \bullet\big) = \big\{\, (q', s') \;\big|\; q' \in \delta(q, \lambda(s)), (s, s') \in R \,\big\}$; and $F'(q,s) = F(q)$.

▶ **Lemma 4.12.** *Let $(q_0, s_0) \bullet (q_1, s_1) \bullet \dots$ be an optimal run of the automaton $\mathcal{A} \times \mathcal{K}$ (that necessarily exists by Lem. 3.2). The path $s_0 s_1 \dots \in \mathrm{path}(\mathcal{K})$ realizes the optimal value of $\mathcal{A}$, that is, $\mathcal{L}(\mathcal{A})\big(\lambda(s_0)\lambda(s_1)\dots\big) = \max_{\xi \in \mathrm{path}(\mathcal{K})} \mathcal{L}(\mathcal{A})\big(\lambda(\xi)\big)$.* ◀

▶ **Theorem 4.13** (optimal scheduling for $\mathrm{LTL}^{\mathrm{disc}}[\mathcal{D}, \emptyset]$)**.** *Assume the setting of Def. 4.2, and that $\mathcal{F} = \emptyset$ (i.e. the formula φ contains no propositional quality operators). Let $(q_0, s_0) \bullet (q_1, s_1) \bullet \dots$ be an optimal run (computed by Lem. 3.2) for the $[0,1]$-acceptance automaton $\mathcal{A}^{\mathrm{na}}_{\varphi,\varepsilon} \times \mathcal{K}$ constructed as in Def. 4.7, Cor. 4.10 and Def. 4.11. Then the path $s_0 s_1 \dots \in \mathrm{path}(\mathcal{K})$ is a solution to the near-optimal scheduling problem (Def. 4.2).*

Moreover, the solution $s_0 s_1 \dots$ can be chosen to be ultimately periodic. ◀

4.2 Our General Algorithm for $\mathrm{LTL}^{\mathrm{disc}}[\mathcal{D}, \mathcal{F}_{\mathrm{mc}}]$

Our general algorithm works in the setting of $\mathrm{LTL}^{\mathrm{disc}}[\mathcal{D}, \mathcal{F}_{\mathrm{mc}}]$—i.e. in the presence of monotone and continuous propositional quality operators like $\oplus$—where threshold model checking is potentially undecidable [3] and therefore the binary-search algorithm (described after Def. 4.2) may not work.

The general algorithm is a (rather straightforward) adaptation of the one we described for $\mathrm{LTL}^{\mathrm{disc}}[\mathcal{D}, \emptyset]$ (§4.1). Here we construct the alternating $[0,1]$-acceptance automaton $\mathcal{A}_{\varphi,\varepsilon}$ *inductively* on the construction on the formula φ:

- When the outermost connective is other than a propositional quality operator, the construction is much like in Def. 4.7.
- When the outermost connective is a propositional quality operator, we rely on Prop. 3.5.

The rest of the algorithm (i.e. the part described in §4.1.3) remains unchanged. An extensive description of the details of the construction is deferred to Appendix A.

▶ **Theorem 4.14** (main theorem, optimal scheduling for $\mathrm{LTL}^{\mathrm{disc}}[\mathcal{D}, \mathcal{F}_{\mathrm{mc}}]$)**.** *In the setting of Def. 4.2, assume that $\mathcal{F} \subseteq \mathcal{F}_{\mathrm{mc}}$ (i.e. all the propositional quality operators in φ are monotone and continuous). Then the near-optimal scheduling problem is decidable.* ◀

4.3 Complexity

The two parameters $\mathcal{D}$ and $\mathcal{F}$ in $\mathrm{LTL}^{\mathrm{disc}}[\mathcal{D}, \mathcal{F}]$—i.e. discounting functions (Def. 2.1) and propositional quality operators (Def. 2.2)—are both relevant to the complexity of our algorithm. Formulating a complexity result is hard when these parameters are left open. We therefore restrict to:

- exponential discounting functions (Def. 2.1), i.e. $\mathcal{D} = \mathcal{D}_{\exp} = \{\exp_\lambda \mid \lambda \in (0,1) \cap \mathbb{Q}\}$, as is done in [3]; and
- the average operator $\oplus$, i.e. $\mathcal{F} = \{\oplus\}$.

We use the definition $|\langle\varphi\rangle|$ of the size of a formula φ, which is from [3]: it reflects the description length of $\lambda \in \mathbb{Q}$ that appears in discounting functions, as well as the length of φ as an expression.

▶ **Proposition 4.15** (size of $\mathcal{A}_{\varphi,\varepsilon}$)**.** *Let φ be an $LTL^{\mathrm{disc}}[\mathcal{D}_{\exp}, \{\oplus\}]$ formula and $\varepsilon \in (0,1) \cap \mathbb{Q}$ be a positive rational number. The size of the state space of the alternating $[0,1]$-acceptance automaton $\mathcal{A}_{\varphi,\varepsilon}$ is singly exponential in $|\langle\varphi\rangle|$ and in the length of the description of ε.* ◀

▶ **Theorem 4.16** (complexity for $\mathrm{LTL}^{\mathrm{disc}}[\mathcal{D}_{\exp}, \{\oplus\}]$)**.** *The near-optimal scheduling problem for $LTL^{\mathrm{disc}}[\mathcal{D}_{\exp}, \{\oplus\}]$ is: in EXPSPACE in $|\langle\varphi\rangle|$ and in the description length of ε; and in NLOGSPACE in the size of $\mathcal{K}$.* ◀

formula φ \ #(states)	$\varepsilon = \frac{1}{10}$		$\varepsilon = \frac{1}{50}$		$\varepsilon = \frac{1}{100}$	
	$\mathcal{A}_{\varphi,\varepsilon}$	$\mathcal{A}^{\mathrm{na}}_{\varphi,\varepsilon}$	$\mathcal{A}_{\varphi,\varepsilon}$	$\mathcal{A}^{\mathrm{na}}_{\varphi,\varepsilon}$	$\mathcal{A}_{\varphi,\varepsilon}$	$\mathcal{A}^{\mathrm{na}}_{\varphi,\varepsilon}$
$\mathsf{F}_{\exp\frac{1}{2}}\, p_1$	5	10	7	14	8	16
$\mathsf{F}_{\exp\frac{99}{100}}\, p_1$	231	462	391	782	460	920
$\mathsf{F}_{\exp\frac{1}{2}}\, \mathsf{G}_{\exp\frac{1}{2}}\, p_1$	15	36	28	85	36	121
$\mathsf{F}_{\exp\frac{1}{2}}\, p_1 \oplus \mathsf{F}_{\exp\frac{1}{2}}\, p_2$	33	128	61	1859	78	7421
$\mathsf{F}_{\exp\frac{1}{2}}\, p_1 \oplus \mathsf{G}_{\exp\frac{1}{2}}\, p_2$	29	272	55	6659	71	32703
$\mathsf{F}_{\exp\frac{3}{5}}\, p_1 \oplus \mathsf{F}_{\exp\frac{3}{5}}\, p_2$	46	477	97	29655	141	timeout (2 min.)
$\mathsf{F}(\mathsf{G}p_1 \oplus \mathsf{F}_{\exp\frac{1}{2}}\, p_2)$	14	19	20	27	23	31

Table 2 Size of the alternating $[0,1]$-acceptance automaton $\mathcal{A}_{\varphi,\varepsilon}$, and $[0,1]$-acceptance automaton $\mathcal{A}_{\varphi,\varepsilon}$

margin ε	#(states of $\mathcal{K}$)	max. outgoing degree of $\mathcal{K}$	time (sec)	space (MB)
$\frac{1}{10}$	100	3	0.085508	5.861
		10	0.114427	9.368
	200	3	0.186989	10.586
		10	0.249392	18.216
$\frac{1}{50}$	100	3	5.928842	199.782
		10	8.108335	405.884
	200	3	10.750703	405.313
		10	18.250345	851.255

Table 3 Time and space consumption of our algorithm for near-optimal scheduling, for the formula $\mathsf{G}_{\exp\frac{1}{2}}\, p_1 \oplus \mathsf{G}_{\exp\frac{1}{2}}\, p_2$ and a randomly generated Kripke structure $\mathcal{K}$. For each choice of the number of states (100 or 200) and of the maximum outgoing degree (3 or 10), we randomly generated 100 instances of $\mathcal{K}$ and the above shows the average

	$\mathsf{G}_{\exp\frac{1}{2}}\, \mathsf{F}p$		$\mathsf{F}_{\exp\frac{1}{2}}\, \mathsf{G}p$	
	time (sec)	space (MB)	time (sec)	space (MB)
our algorithm (§4.1)	18.918600	897.111	0.019800	4.629
binary search	0.047200	5.140390	0.069500	5.567

Table 4 (Comparison with binary search, in absence of $\oplus$) Time and space consumption for near-optimal scheduling, for the margin $\varepsilon = \frac{1}{100}$ and a randomly generated Kripke structure $\mathcal{K}$ (500 states, max. outgoing degree 10, average over 100 instances)

In case of absence of propositional quality operators (i.e. $\mathrm{LTL}^{\mathrm{disc}}[\mathcal{D}_{\exp}, \emptyset]$), we can further optimize the complexity by using a heuristic and avoiding the exponential blowup from $\mathcal{A}_{\varphi,\varepsilon}$ to $\mathcal{A}^{\mathrm{na}}_{\varphi,\varepsilon}$. This yields the following complexity result, which is also achievable by the binary-search algorithm.

▶ **Theorem 4.17** (complexity for $\mathrm{LTL}^{\mathrm{disc}}[\mathcal{D}_{\exp}, \emptyset]$)**.** *The near-optimal scheduling problem for $\mathrm{LTL}^{\mathrm{disc}}[\mathcal{D}_{\exp}, \emptyset]$ is: in PSPACE in $|\langle\varphi\rangle|$ and in the description length of ε; and in NLOGSPACE in the size of $\mathcal{K}$.* ◀

5 Experiments

We implemented our algorithm in §4 that solves the near-optimal scheduling for $\mathrm{LTL}^{\mathrm{disc}}[\mathcal{D}_{\exp}, \{\oplus\}]$. The implementation is in OCaml. The following experiments were on a MacBook Pro laptop with a Core i5 processor (2.7 GHz) and 16 GB RAM. In Table 2, for each choice of φ and ε, we show the size of the alternating automaton $\mathcal{A}_{\varphi,\varepsilon}$, and the non-alternating $\mathcal{A}^{\mathrm{na}}_{\varphi,\varepsilon}$ that

results from $\mathcal{A}_{\varphi,\varepsilon}$. The first three rows have no $\oplus$, in which case the implementation scales well for bigger bases (i.e. discount functions that decrease more slowly). We observe that presence of $\oplus$ incurs substantial computational costs: the small increase of bases from $\frac{1}{2}$ (the fourth row) to $\frac{3}{5}$ (the sixth row) makes $\mathcal{A}_{\varphi,\varepsilon}$ much bigger, resulting in one timeout. This is as expected, however: $\oplus$ makes other problems harder too, such as model checking (undecidable).

In Table 3 we fix a formula $\varphi = \mathsf{G}_{\exp\frac{1}{2}}\, p_1 \oplus \mathsf{G}_{\exp\frac{1}{2}}\, p_2$ and measure time and space consumption, for various choices of a margin ε and a Kripke structure $\mathcal{K}$. Kripke structures $\mathcal{K}$ were randomly generated: we first set the number of states (100 or 200) and the maximum outgoing degree of $\mathcal{K}$ (3 or 10); for each state we fixed its outgoing degree, from the uniform distribution from 1 to the maximum (that we had already fixed); and then, for each outgoing edge, its target state is chosen from the uniform distribution over the set of states. We observe that time and space consumption grows significantly as the problem becomes more difficult. However, for problem instances of a considerable size we still see manageable costs: a margin $\varepsilon = \frac{1}{50}$ (2%) is fairly small, and a Kripke structure $\mathcal{K}$ with 200 states is likely to be capable of modeling many communication protocols.

In Table 4, for reference, we compare our algorithm in §4.1 with the binary-search algorithm that exploits the model-checking algorithm in [3] (we also implemented the latter). We emphasize again that the latter does not work in presence of $\oplus$. Our experience shows that the binary-search algorithm can in some cases be faster by a magnitude (e.g. for the first formula here), but not always (for the second formula our algorithm is a few times faster).

Those experimental results indicate that, although presence of the average operator $\oplus$ incurs significant computational cost (as expected), automata-based optimal scheduling for $\mathrm{LTL}^{\mathrm{disc}}[\mathcal{D}_{\exp}, \{\oplus\}]$ is potentially a viable approach. It is not that our algorithm scales up to huge problem instances, but systems of hundreds of states can be handled without difficulties. Identification of concrete real-world challenges, and enhancement of the tool's efficiency to match up to them, is an important direction of future work.

6 Conclusions and Future Work

For the quantitative logic $\mathrm{LTL}^{\mathrm{disc}}[\mathcal{D}_{\exp}, \mathcal{F}]$ with future discounting [3], we formulated a natural problem of synthesizing near-optimal schedulers, and presented an algorithm. The latter relies on: the existing idea of *event horizon* exploited in [3] for the threshold model checking problem, as well as a supposedly widely-applicable technique of translation to $[0,1]$-acceptance automata and a lasso-style optimal value algorithm for them.

Here are several directions of future work.

Controller Synthesis for Open Systems We note that the current results are focused on *closed* systems. For *open* or *reactive* systems (like a server that responds to requests that come from the environment) we would wish to synthesize a *controller*—formally a *strategy* or a *transducer*—that achieves a near-optimal performance.

An envisaged workflow, following the one in [22], is as follows. We will use the same automaton $\mathcal{A}_{\varphi,\varepsilon}$ (Def. 4.7). It is then: 1) determinized, 2) transformed into a tree automaton that accepts the desired strategies, and 3) the optimal value of the tree automaton is checked, much like in Lem. 3.2. While the step 2) will be straightforward, the steps 1) and 3) (namely: determinization of $[0,1]$-acceptance automata, and the optimal value problem for "$[0,1]$-acceptance Rabin automata") are yet to be investigated. Another possible workflow is by an adaptation of the Safraless algorithm [16].

Probabilistic Systems and LTL$^{\text{disc}}[\mathcal{D}_{\text{exp}}, \mathcal{F}]$ Here and in [3] the system model is a Kripke structure that is nondeterministic. Adding probabilistic branching will gives us a set of new problems to be solved: for Markov chains the threshold model-checking problem can be formulated; for Markov decision processes, we have both the threshold model-checking problem and the near-optimal scheduling problem. Furthermore, another axis of variation is given by whether we consider the expected value or the worst-case value. In the latter case we would wish to exclude truth values that arise with probability 0. All these variations have important applications in various areas.

Acknowledgments

Thanks are due to Shaull Almagor, Shuichi Hirahara, and the anonymous referees, for useful discussions and comments. The authors are supported by Grants-in-Aid No. 24680001, 15KT0012 and 15K11984, JSPS.

References

1 Yasmina Abdeddaïm, Eugene Asarin, and Oded Maler. Scheduling with timed automata. *Theor. Comput. Sci.*, 354(2):272–300, 2006.

2 Shaull Almagor, Udi Boker, and Orna Kupferman. Formalizing and reasoning about quality. In Fedor V. Fomin, Rusins Freivalds, Marta Z. Kwiatkowska, and David Peleg, editors, *Automata, Languages, and Programming - 40th International Colloquium, ICALP 2013, Riga, Latvia, July 8-12, 2013, Proceedings, Part II*, volume 7966 of *Lecture Notes in Computer Science*, pages 15–27. Springer, 2013.

3 Shaull Almagor, Udi Boker, and Orna Kupferman. Discounting in LTL. In Erika Ábrahám and Klaus Havelund, editors, *Tools and Algorithms for the Construction and Analysis of Systems - 20th International Conference, TACAS 2014, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2014, Grenoble, France, April 5-13, 2014. Proceedings*, volume 8413 of *Lecture Notes in Computer Science*, pages 424–439. Springer, 2014.

4 Shaull Almagor, Udi Boker, and Orna Kupferman. Formalizing and reasoning about quality. Extended version of [2], preprint (private communication), 2014.

5 Christel Baier, Clemens Dubslaff, and Sascha Klüppelholz. Trade-off analysis meets probabilistic model checking. In Thomas A. Henzinger and Dale Miller, editors, *Joint Meeting of the Twenty-Third EACSL Annual Conference on Computer Science Logic (CSL) and the Twenty-Ninth Annual ACM/IEEE Symposium on Logic in Computer Science (LICS), CSL-LICS '14, Vienna, Austria, July 14 - 18, 2014*, page 1. ACM, 2014.

6 Roderick Bloem, Krishnendu Chatterjee, Thomas A. Henzinger, and Barbara Jobstmann. Better quality in synthesis through quantitative objectives. In Ahmed Bouajjani and Oded Maler, editors, *Computer Aided Verification, 21st International Conference, CAV 2009, Grenoble, France, June 26 - July 2, 2009. Proceedings*, volume 5643 of *Lecture Notes in Computer Science*, pages 140–156. Springer, 2009.

7 Patricia Bouyer, Nicolas Markey, and Raj Mohan Matteplackel. Averaging in LTL. In Paolo Baldan and Daniele Gorla, editors, *CONCUR 2014 - Concurrency Theory - 25th International Conference, CONCUR 2014, Rome, Italy, September 2-5, 2014. Proceedings*, volume 8704 of *Lecture Notes in Computer Science*, pages 266–280. Springer, 2014.

8 Pavol Cerný, Krishnendu Chatterjee, Thomas A. Henzinger, Arjun Radhakrishna, and Rohit Singh. Quantitative synthesis for concurrent programs. In Ganesh Gopalakrishnan and Shaz Qadeer, editors, *Computer Aided Verification - 23rd International Conference, CAV 2011, Snowbird, UT, USA, July 14-20, 2011. Proceedings*, volume 6806 of *Lecture Notes in Computer Science*, pages 243–259. Springer, 2011.

9 Krishnendu Chatterjee, Laurent Doyen, and Thomas A. Henzinger. Expressiveness and closure properties for quantitative languages. *Logical Methods in Computer Science*, 6(3), 2010.

10 Krishnendu Chatterjee, Thomas A. Henzinger, and Marcin Jurdzinski. Mean-payoff parity games. In *20th IEEE Symposium on Logic in Computer Science (LICS 2005), 26-29 June 2005, Chicago, IL, USA, Proceedings*, pages 178–187. IEEE Computer Society, 2005.

11 Ling Cheung, Mariëlle Stoelinga, and Frits W. Vaandrager. A testing scenario for probabilistic processes. *J. ACM*, 54(6), 2007.

12 Luca de Alfaro, Thomas A. Henzinger, and Rupak Majumdar. Discounting the future in systems theory. In Jos C. M. Baeten, Jan Karel Lenstra, Joachim Parrow, and Gerhard J. Woeginger, editors, *Automata, Languages and Programming, 30th International Colloquium, ICALP 2003, Eindhoven, The Netherlands, June 30 - July 4, 2003. Proceedings*, volume 2719 of *Lecture Notes in Computer Science*, pages 1022–1037. Springer, 2003.

13 Manfred Droste and Ulrike Püschmann. On weighted Büchi automata with order-complete weights. *IJAC*, 17(2):235–260, 2007.

14 Marco Faella, Axel Legay, and Mariëlle Stoelinga. Model checking quantitative linear time logic. *Electr. Notes Theor. Comput. Sci.*, 220(3):61–77, 2008.

15 Hans Hansson and Bengt Jonsson. A logic for reasoning about time and reliability. *Formal Asp. Comput.*, 6(5):512–535, 1994.

16 Orna Kupferman, Nir Piterman, and Moshe Y. Vardi. Safraless compositional synthesis. In Thomas Ball and Robert B. Jones, editors, *Computer Aided Verification, 18th International Conference, CAV 2006, Seattle, WA, USA, August 17-20, 2006, Proceedings*, volume 4144 of *Lecture Notes in Computer Science*, pages 31–44. Springer, 2006.

17 Orna Kupferman, Moshe Y. Vardi, and Pierre Wolper. An automata-theoretic approach to branching-time model checking. *J. ACM*, 47(2):312–360, March 2000.

18 Satoru Miyano and Takeshi Hayashi. Alternating finite automata on omega-words. *Theor. Comput. Sci.*, 32:321–330, 1984.

19 Amir Pnueli and Roni Rosner. On the synthesis of a reactive module. In *Conference Record of the Sixteenth Annual ACM Symposium on Principles of Programming Languages, Austin, Texas, USA, January 11-13, 1989*, pages 179–190. ACM Press, 1989.

20 George Rahonis. Infinite fuzzy computations. *Fuzzy Sets and Systems*, 153(2):275–288, 2005.

21 R. J. van Glabbeek. The linear time–branching time spectrum I; the semantics of concrete, sequential processes. In J. A. Bergstra, A. Ponse, and S. A. Smolka, editors, *Handbook of Process Algebra*, chapter 1, pages 3–99. Elsevier, 2001.

22 Moshe Y. Vardi. An automata-theoretic approach to linear temporal logic. In *Logics for Concurrency: Structure versus Automata, volume 1043 of Lecture Notes in Computer Science*, pages 238–266. Springer-Verlag, 1996.

A Our General Algorithm for $\mathrm{LTL}^{\mathrm{disc}}[\mathcal{D}, \mathcal{F}_{\mathrm{mc}}]$, Further Details

In this section, we extend §4.2 and describe details of the construction of $\mathcal{A}_{\varphi,\varepsilon}$ for a formula φ of $\mathrm{LTL}^{\mathrm{disc}}[\mathcal{D}, \mathcal{F}_{\mathrm{mc}}]$. We inductively construct an alternating $[0,1]$-acceptance automaton $\mathcal{A}^{\vec{d}}_{\varphi,\varepsilon}$—that is also parametrized by a discount sequence $\vec{d}$. Then the automaton $\mathcal{A}_{\varphi,\varepsilon}$ is defined by $\mathcal{A}^{\langle 1\rangle}_{\varphi,\varepsilon}$ (for the sequence $\langle 1\rangle$ of length one).

▶ **Lemma A.1.** *Let φ be an $LTL^{\mathrm{disc}}[\mathcal{D}, \mathcal{F}_{\mathrm{mc}}]$ formula, $\varepsilon \in (0,1)$ be a positive real number, and $\vec{d} = d_1 d_2 \dots d_n$ be a discount sequence (Def. 4.4). There exists an alternating $[0,1]$-acceptance automaton $\mathcal{A}^{\vec{d}}_{\varphi,\varepsilon}$ such that, for each computation $\pi \in (\mathcal{P}(AP))^\omega$,*

$$\big(\vec{d} \boxtimes [\![\pi, \varphi]\!]\big) - \varepsilon \;\le\; \mathcal{L}(\mathcal{A}^{\vec{d}}_{\varphi,\varepsilon})(\pi) \;\le\; \vec{d} \boxtimes [\![\pi, \varphi]\!] \;. \tag{15}$$

Proof. The proof is inductive on the construction of φ.[1] In this proof we assume without loss of generality that an alternating $[0,1]$-acceptance automaton has exactly one initial state, and consequently, the initial state of $\mathcal{A}^{\vec{d}}_{\varphi,\varepsilon}$ shall be denoted by $q^{\vec{d}}_{\varphi,\varepsilon}$. For the case where the outermost connective of φ is other than a propositional quality operator, we only describe the construction of $\mathcal{A}^{\vec{d}}_{\varphi,\varepsilon}$. The correctness of this automaton can be proved in a similar way to the proof of Lem. 4.9: recall that Lem. 4.9 is also proved by induction on the construction of a formula.

Suppose that $\varphi = \mathtt{True}$. We define $\mathcal{A}^{\vec{d}}_{\mathtt{True},\varepsilon} = (\mathcal{P}(AP), \{q^{\vec{d}}_{\mathtt{True},\varepsilon}\}, \{q^{\vec{d}}_{\mathtt{True},\varepsilon}\}, \delta, F)$ where $\delta(q^{\vec{d}}_{\mathtt{True},\varepsilon}, \sigma) = \vec{d} \boxtimes 1$ and $F(q^{\vec{d}}_{\mathtt{True},\varepsilon}) = 0$.

Suppose that $\varphi = p \in AP$. We define $\mathcal{A}^{\vec{d}}_{p,\varepsilon} = (\mathcal{P}(AP), \{q^{\vec{d}}_{p,\varepsilon}\}, \{q^{\vec{d}}_{p,\varepsilon}\}, \delta, F)$ where

$$\delta(q^{\vec{d}}_{p,\varepsilon}, \sigma) = \begin{cases} \vec{d} \boxtimes 1 & \text{if } p \in \sigma \\ \vec{d} \boxtimes 0 & \text{otherwise} \end{cases} \qquad \text{and} \qquad F(q^{\vec{d}}_{p,\varepsilon}) = 0 \;.$$

Suppose that $\varphi = \varphi_1 \wedge \varphi_2$ and that $|\vec{d}|$ is odd. By the induction hypothesis, for each of $i \in \{1,2\}$, there exists $\mathcal{A}^{\vec{d}}_{\varphi_i,\varepsilon} = (\mathcal{P}(AP), Q_i, \{q^{\vec{d}}_{\varphi_i,\varepsilon}\}, \delta_i, F_i)$ that satisfies the postulated condition. Then we define $\mathcal{A}^{\vec{d}}_{\varphi,\varepsilon} = (\mathcal{P}(AP), Q, \{q^{\vec{d}}_{\varphi,\varepsilon}\}, \delta, F)$ as follows. Its state space Q is $\{q^{\vec{d}}_{\varphi,\varepsilon}\} \cup Q_1 \cup Q_2$. The transition function δ is

$$\delta(q,\sigma) = \begin{cases} \delta_1(q^{\vec{d}}_{\varphi_1,\varepsilon}, \sigma) \wedge \delta_2(q^{\vec{d}}_{\varphi_2,\varepsilon}, \sigma) & \text{if } q = q^{\vec{d}}_{\varphi,\varepsilon} \\ \delta_i(q,\sigma) & \text{if } q \in Q_i. \end{cases}$$

The acceptance function F is

$$F(q) = \begin{cases} 0 & \text{if } q = q^{\vec{d}}_{\varphi,\varepsilon} \\ F_i(q) & \text{if } q \in Q_i. \end{cases}$$

Suppose that $\varphi = \varphi_1 \wedge \varphi_2$ and that $|\vec{d}|$ is even. By the induction hypothesis, for each of $i \in \{1,2\}$, there exists $\mathcal{A}^{\vec{d}}_{\varphi_i,\varepsilon} = (\mathcal{P}(AP), Q_i, \{q^{\vec{d}}_{\varphi_i,\varepsilon}\}, \delta_i, F_i)$ that satisfies the postulated condition. Then we define $\mathcal{A}^{\vec{d}}_{\varphi,\varepsilon} = (\mathcal{P}(AP), Q, \{q^{\vec{d}}_{\varphi,\varepsilon}\}, \delta, F)$ as follows. Its state space Q is $\{q^{\vec{d}}_{\varphi,\varepsilon}\} \cup Q_1 \cup Q_2$. The transition function δ is

$$\delta(q,\sigma) = \begin{cases} \delta_1(q^{\vec{d}}_{\varphi_1,\varepsilon}, \sigma) \vee \delta_2(q^{\vec{d}}_{\varphi_2,\varepsilon}, \sigma) & \text{if } q = q^{\vec{d}}_{\varphi,\varepsilon} \\ \delta_i(q,\sigma) & \text{if } q \in Q_i. \end{cases}$$

[1] To be precise, we have two nested induction: the outer one is with respect to the number of propositional quality operators occurring in φ; and the inner one is with respect to the size of a formula φ.

The acceptance function F is

$$F(q) = \begin{cases} 0 & \text{if } q = q^{\vec{d}}_{\varphi,\varepsilon} \\ F_i(q) & \text{if } q \in Q_i. \end{cases}$$

Suppose that $\varphi = \neg\varphi'$. By the induction hypothesis, there exists $\mathcal{A}^{\vec{d}:1}_{\varphi',\varepsilon}$ that satisfies the postulated condition. Let $\mathcal{A}^{\vec{d}}_{\varphi,\varepsilon} = \mathcal{A}^{\vec{d}:1}_{\varphi',\varepsilon}$.

Suppose that $\varphi = \mathsf{X}\varphi'$. By the induction hypothesis, there exists $\mathcal{A}^{\vec{d}}_{\varphi',\varepsilon} = (\mathcal{P}(AP), Q', \{q^{\vec{d}}_{\varphi',\varepsilon}\}, \delta', F')$ that satisfies the postulated condition. Then we define $\mathcal{A}^{\vec{d}}_{\varphi,\varepsilon} = (\mathcal{P}(AP), Q, \{q^{\vec{d}}_{\varphi,\varepsilon}\}, \delta, F)$ as follows. Its state space Q is $\{q^{\vec{d}}_{\varphi,\varepsilon}\} \cup Q'$. The transition function δ is

$$\delta(q,\sigma) = \begin{cases} q^{\vec{d}}_{\varphi',\varepsilon} & \text{if } q = q^{\vec{d}}_{\varphi,\varepsilon} \\ \delta'(q,\sigma) & \text{otherwise.} \end{cases}$$

The acceptance function F is

$$F(q) = \begin{cases} 0 & \text{if } q = q^{\vec{d}}_{\varphi,\varepsilon} \\ F'(q) & \text{otherwise.} \end{cases}$$

Suppose that $\varphi = \varphi_1 \mathsf{U} \varphi_2$ and that $|\vec{d}|$ is odd. By the induction hypothesis, for each of $i \in \{1,2\}$, there exists $\mathcal{A}^{\vec{d}}_{\varphi_i,\varepsilon} = (\mathcal{P}(AP), Q_i, \{q^{\vec{d}}_{\varphi_i,\varepsilon}\}, \delta_i, F_i)$ that satisfies the postulated condition. Then we define $\mathcal{A}^{\vec{d}}_{\varphi,\varepsilon} = (\mathcal{P}(AP), Q, \{q^{\vec{d}}_{\varphi,\varepsilon}\}, \delta, F)$ as follows. Its state space Q is $\{q^{\vec{d}}_{\varphi,\varepsilon}\} \cup Q_1 \cup Q_2$. The transition function δ is

$$\delta(q,\sigma) = \begin{cases} \delta_2(q^{\vec{d}}_{\varphi_2,\varepsilon},\sigma) \vee (\delta_1(q^{\vec{d}}_{\varphi_1,\varepsilon},\sigma) \wedge q^{\vec{d}}_{\varphi,\varepsilon}) & \text{if } q = q^{\vec{d}}_{\varphi,\varepsilon} \\ \delta_i(q,\sigma) & \text{if } q \in Q_i. \end{cases}$$

The acceptance function F is

$$F(q) = \begin{cases} 0 & \text{if } q = q^{\vec{d}}_{\varphi,\varepsilon} \\ F_i(q) & \text{if } q \in Q_i. \end{cases}$$

Suppose that $\varphi = \varphi_1 \mathsf{U} \varphi_2$ and that $|\vec{d}|$ is even. By the induction hypothesis, for each of $i \in \{1,2\}$, there exists $\mathcal{A}^{\vec{d}}_{\varphi_i,\varepsilon} = (\mathcal{P}(AP), Q_i, \{q^{\vec{d}}_{\varphi_i,\varepsilon}\}, \delta_i, F_i)$ that satisfies the postulated condition. Then we define $\mathcal{A}^{\vec{d}}_{\varphi,\varepsilon} = (\mathcal{P}(AP), Q, \{q^{\vec{d}}_{\varphi,\varepsilon}\}, \delta, F)$ as follows. Its state space Q is $\{q^{\vec{d}}_{\varphi,\varepsilon}\} \cup Q_1 \cup Q_2$. The transition function δ is

$$\delta(q,\sigma) = \begin{cases} \delta_2(q^{\vec{d}}_{\varphi_2,\varepsilon},\sigma) \wedge (\delta_1(q^{\vec{d}}_{\varphi_1,\varepsilon},\sigma) \vee q^{\vec{d}}_{\varphi,\varepsilon}) & \text{if } q = q^{\vec{d}}_{\varphi,\varepsilon} \\ \delta_i(q,\sigma) & \text{if } q \in Q_i. \end{cases}$$

The acceptance function F is

$$F(q) = \begin{cases} 1 & \text{if } q = q^{\vec{d}}_{\varphi,\varepsilon} \\ F_i(q) & \text{if } q \in Q_i. \end{cases}$$

Suppose that $\varphi = \varphi_1 \mathsf{U}_{\eta^{+k}} \varphi_2$ and that $|\vec{d}|$ is odd. Since $\lim_{i\to\infty} \eta(i) = 0$, there exists a natural number $k_{\max} \in \mathbb{N}$ such that $\eta(k_{\max}) \cdot \prod_{i=1}^{n} d_i \leq \varepsilon$ (i.e. $k_{\max}$ is beyond the event horizon). We construct $\mathcal{A}^{\vec{d}}_{\varphi,\varepsilon}$ by induction on k backwards, that is, starting from $k = k_{\max}$

and decrementing k one by one until $k = 0$. If $\eta(k) \cdot \prod_{i=1}^{n} d_i \leq \varepsilon$, we define $\mathcal{A}^{\vec{d}}_{\varphi,\varepsilon} = (\mathcal{P}(AP), \{q^{\vec{d}}_{\varphi,\varepsilon}\}, \{q^{\vec{d}}_{\varphi,\varepsilon}\}, \delta, F)$ where $\delta(q^{\vec{d}}_{\varphi,\varepsilon}, \sigma) = \vec{d} \boxtimes 0$ and $F(q^{\vec{d}}_{\varphi,\varepsilon}) = 0$. Otherwise, we define $\mathcal{A}^{\vec{d}}_{\varphi,\varepsilon} = (\mathcal{P}(AP), Q, \{q^{\vec{d}}_{\varphi,\varepsilon}\}, \delta, F)$ as follows. By the induction hypothesis, for each of $i \in \{1, 2\}$, there exists $\mathcal{A}^{\vec{d} \odot \eta(k)}_{\varphi_i,\varepsilon} = (\mathcal{P}(AP), Q_i, \{q^{\vec{d} \odot \eta(k)}_{\varphi_i,\varepsilon}\}, \delta_i, F_i)$ that satisfies the postulated condition. Moreover, there exists $\mathcal{A}^{\vec{d}}_{\varphi_1 \mathsf{U}_{\eta^{k+1}} \varphi_2,\varepsilon} = (\mathcal{P}(AP), Q_3, \{q^{\vec{d}}_{\varphi_1 \mathsf{U}_{\eta^{k+1}} \varphi_2,\varepsilon}\}, \delta_3, F_3)$. We define the state space Q of $\mathcal{A}^{\vec{d}}_{\varphi,\varepsilon}$ by $\{q^{\vec{d}}_{\varphi,\varepsilon}\} \cup Q_1 \cup Q_2 \cup Q_3$. The transition function δ is

$$\delta(q, \sigma) = \begin{cases} \delta_2(q^{\vec{d} \odot \eta(k)}_{\varphi_2,\varepsilon}, \sigma) \vee (\delta_1(q^{\vec{d} \odot \eta(k)}_{\varphi_1,\varepsilon}, \sigma) \wedge q^{\vec{d}}_{\varphi_1 \mathsf{U}_{\eta^{k+1}} \varphi_2,\varepsilon}) & \text{if } q = q^{\vec{d}}_{\varphi,\varepsilon} \\ \delta_i(q, \sigma) & \text{if } q \in Q_i. \end{cases}$$

The acceptance function F is

$$F(q) = \begin{cases} 0 & \text{if } q = q^{\vec{d}}_{\varphi,\varepsilon} \\ F_i(q) & \text{if } q \in Q_i. \end{cases}$$

Suppose that $\varphi = \varphi_1 \, \mathsf{U}_{\eta^{+k}} \, \varphi_2$ and that $|\vec{d}|$ is even. Similarly to the case where $|\vec{d}|$ is odd, we construct $\mathcal{A}^{\vec{d}}_{\varphi,\varepsilon}$ by induction on k backwards. If $\eta(k) \cdot \prod_{i=1}^{n} d_i \leq \varepsilon$, we define $\mathcal{A}^{\vec{d}}_{\varphi,\varepsilon} = (\mathcal{P}(AP), \{q^{\vec{d}}_{\varphi,\varepsilon}\}, \{q^{\vec{d}}_{\varphi,\varepsilon}\}, \delta, F)$ where $\delta(q^{\vec{d}}_{\varphi,\varepsilon}, \sigma) = \vec{d} \boxtimes \eta(k)$ and $F(q^{\vec{d}}_{\varphi,\varepsilon}) = 0$. Otherwise, we define $\mathcal{A}^{\vec{d}}_{\varphi,\varepsilon} = (\mathcal{P}(AP), Q, \{q^{\vec{d}}_{\varphi,\varepsilon}\}, \delta, F)$ as follows. By the induction hypothesis, for each of $i \in \{1, 2\}$, there exists $\mathcal{A}^{\vec{d} \odot \eta(k)}_{\varphi_i,\varepsilon} = (\mathcal{P}(AP), Q_i, \{q^{\vec{d} \odot \eta(k)}_{\varphi_i,\varepsilon}\}, \delta_i, F_i)$ that satisfies the postulated condition. Moreover, there exists $\mathcal{A}^{\vec{d}}_{\varphi_1 \mathsf{U}_{\eta^{k+1}} \varphi_2,\varepsilon} = (\mathcal{P}(AP), Q_3, \{q^{\vec{d}}_{\varphi_1 \mathsf{U}_{\eta^{k+1}} \varphi_2,\varepsilon}\}, \delta_3, F_3)$. We define the state space Q of $\mathcal{A}^{\vec{d}}_{\varphi,\varepsilon}$ by $\{q^{\vec{d}}_{\varphi,\varepsilon}\} \cup Q_1 \cup Q_2 \cup Q_3$. The transition function δ is

$$\delta(q, \sigma) = \begin{cases} \delta_2(q^{\vec{d} \odot \eta(k)}_{\varphi_2,\varepsilon}, \sigma) \wedge (\delta_1(q^{\vec{d} \odot \eta(k)}_{\varphi_1,\varepsilon}, \sigma) \vee q^{\vec{d}}_{\varphi_1 \mathsf{U}_{\eta^{k+1}} \varphi_2,\varepsilon}) & \text{if } q = q^{\vec{d}}_{\varphi,\varepsilon} \\ \delta_i(q, \sigma) & \text{if } q \in Q_i. \end{cases}$$

The acceptance function F is

$$F(q) = \begin{cases} 0 & \text{if } q = q^{\vec{d}}_{\varphi,\varepsilon} \\ F_i(q) & \text{if } q \in Q_i. \end{cases}$$

Suppose that $\varphi = f(\varphi_1, \ldots, \varphi_k)$ where $f \in \mathcal{F}_{\mathrm{mc}}$ and that $|\vec{d}|$ is odd. Since f is continuous and its domain $[0, 1]^k$ is bounded and closed in the Euclidean space $\mathbb{R}^k$, this function f is uniformly continuous by the Heine–Cantor theorem. By the monotonicity and the uniform continuity, there exists $\varepsilon' \in (0, 1)$ such that, for each $\mathbf{x} = (x_1, \ldots, x_k) \in [0, 1]^k$,

$$f(x_1 \dot{-} \varepsilon', \ldots, x_k \dot{-} \varepsilon') \geq f(\mathbf{x}) - \varepsilon/(d_1 \cdot d_2 \cdot \cdots \cdot d_n) \ , \tag{16}$$

where $a \dot{-} b$ is defined by $\max\{a - b, 0\}$. By the induction hypothesis, there exist $\mathcal{A}^{\langle 1 \rangle}_{\varphi_1,\varepsilon'}, \ldots, \mathcal{A}^{\langle 1 \rangle}_{\varphi_k,\varepsilon'}$ such that, for $i \in \{1, \ldots, k\}$,

$$[\![\pi, \varphi_i]\!] - \varepsilon' \ \leq \ \mathcal{L}(\mathcal{A}^{\langle 1 \rangle}_{\varphi_i,\varepsilon'})(\pi) \ \leq \ [\![\pi, \varphi_i]\!]$$

for each $\pi \in (\mathcal{P}(AP))^\omega$. Since $|\vec{d}|$ is odd, the function $g \colon [0, 1]^k \to [0, 1]$ defined by $g(\mathbf{x}) = \vec{d} \boxtimes f(\mathbf{x})$ is monotone in $\mathbf{x}$. Since the class of languages of alternating $[0, 1]$-acceptance automata and that of $[0, 1]$-acceptance automata are the same by Prop. 3.4, the closure property in

Prop. 3.5 remains true even if $[0,1]$-acceptance automata are replaced by alternating $[0,1]$-acceptance automata. Hence, there exists $g(\mathcal{A}^{\langle 1\rangle}_{\varphi_1,\varepsilon'},\dots,\mathcal{A}^{\langle 1\rangle}_{\varphi_k,\varepsilon'})$ defined in Prop. 3.5. By (16) and the definition (7) of the operator $\boxtimes$, we have

$$g(x_1 \dot{-} \varepsilon',\dots,x_1 \dot{-} \varepsilon') \geq g(\mathbf{x}) - \varepsilon = \vec{d} \boxtimes f(\mathbf{x}) - \varepsilon \ .$$

Hence, if we define $\mathcal{A}^{\vec{d}}_{\varphi,\varepsilon}$ by $g(\mathcal{A}^{\langle 1\rangle}_{\varphi_1,\varepsilon'},\dots,\mathcal{A}^{\langle 1\rangle}_{\varphi_k,\varepsilon'})$, it satisfies the postulated condition.

Suppose that $\varphi = f(\varphi_1,\dots,\varphi_k)$ where $f \in \mathcal{F}_{\mathrm{mc}}$ and that $|\vec{d}|$ is even. Let $\vec{d'} = d_1 d_2 \dots d_{n-1}$ be a prefix of $\vec{d}$. Then we have $\vec{d} \boxtimes v = \vec{d'} \boxtimes (1 - d_n \cdot v)$. We define a function $(d_n \cdot f)^*\colon [0,1]^k \to [0,1]$ by $(d_n \cdot f)^*(x_1,\dots,x_k) = 1 - d_n \cdot f(1-x_1,\dots,1-x_k)$. Let $\varphi' = (d_n \cdot f)^*(\neg\varphi_1,\dots,\neg\varphi_k)$. It is obvious that $(d_n \cdot f)^* \in \mathcal{F}_{\mathrm{mc}}$. Moreover, we have $\vec{d} \boxtimes [\![\pi,\varphi]\!] = \vec{d'} \boxtimes [\![\pi,\varphi']\!]$ for each $\pi \in (\mathcal{P}(AP))^\omega$, and $\vec{d'}$ is odd. Therefore there exists $\mathcal{A}^{\vec{d'}}_{\varphi',\varepsilon}$ because of the previous case (i.e. when $|\vec{d}|$ is odd),[2] and we take this as $\mathcal{A}^{\vec{d}}_{\varphi,\varepsilon}$. Then $\mathcal{A}^{\vec{d}}_{\varphi,\varepsilon}$ satisfies the postulated condition. ◀

Once $\mathcal{A}_{\varphi,\varepsilon}$ is constructed, the procedure described in §4.1.3 works regardless of the presence of propositional quality operators.

B Omitted Proofs

B.1 Proof of Prop. 3.4

Proof. We first describe the formal construction; intuitions follow shortly.

Without loss of generality, we can assume that a positive Boolean formula $\delta(q,a)$ is a disjunctive normal form; therefore the transition function is of the type $\delta : Q \times \Sigma \to \mathcal{P}(\mathcal{P}(Q \cup [0,1]))$. More concretely, for each $q \in Q$ and $a \in \Sigma$, the formula $\delta(q,a)$ is a disjunction of formulas of the form

$$(q_1 \wedge \dots \wedge q_k) \wedge (v_1 \wedge \dots \wedge v_l)$$

where $q_j \in Q$ and $v_j \in [0,1]$ are atomic propositions (we changed their order suitably). Moreover, since the conjunction $v_1 \wedge \dots \wedge v_l$ is equivalent to a single atomic proposition $\min\{v_1,\dots,v_l\}$, we assume that any disjunct of the DNF formula $\delta(q,a)$ is of the form

$$(q_1 \wedge \dots \wedge q_k) \wedge v \ .$$

Let $V_Q = \{F(q) \mid q \in Q\}$ be the set of acceptance values that occur in $\mathcal{A}$, and V_δ be the set of values from $[0,1]$ (i.e. atomic propositions from $[0,1]$) that occur in the transition function δ, that is,

$$V_\delta \ = \ \bigcup_{q\in Q, a\in\Sigma} \big\{\, v \,\big|\, \big((q_1 \wedge \dots \wedge q_k) \wedge v\big) \in \delta(q,a) \,\big\} \ .$$

We define $\mathcal{A}' = (\Sigma, Q', I', \delta', F')$ as follows.

$$\begin{aligned} Q' \ &= \ \mathcal{P}(Q \times V_Q) \times V_\delta \times \{\mathtt{ff},\mathtt{tt}\} \ , \\ I' \ &= \ \big\{\, \big(\{(q_0, F(q_0))\}, 1, \mathtt{ff}\big) \,\big|\, q_0 \in I \,\big\} \ , \\ F'(Y,v,b) \ &= \ \begin{cases} \min\big\{\, v, \min\{v' \mid \exists q \in Q.\, (q,v') \in Y\}\,\big\} & \text{if } b = \mathtt{tt} \\ 0 & \text{otherwise.} \end{cases} \end{aligned}$$

[2] Recall that we are currently running two nested induction, with the outer one being with respect to the number of propositional quality operators.

The transition function δ' is defined as follows. Let $\widetilde{q} = \big(\{\!\{ (q^1, v^1), \dots, (q^n, v^n) \}\!\}, v, b \big)$ be a state in Q', and $a \in \Sigma$. Then $\delta'(\widetilde{q}, a)$ is defined, in case $b = \mathtt{ff}$, by:

$$\left\{ \left(\begin{Bmatrix} \big(q_1^1,\, \max\{v^1, F(q_1^1)\}\big)\,, \;\dots,\; \big(q_{l_1}^1,\, \max\{v^1, F(q_{l_1}^1)\}\big), \\ \vdots \\ \big(q_1^n,\, \max\{v^n, F(q_1^n)\}\big)\,, \;\dots,\; \big(q_{l_n}^n,\, \max\{v^n, F(q_{l_n}^n)\}\big) \end{Bmatrix}, \begin{array}{c} \min\{v, u^1, \dots, u^n\}\,, \\ b' \end{array} \right) \;\middle|\; \big((q_1^i \wedge \dots \wedge q_{l_i}^i) \wedge u^i \big) \in \delta(q^i, a)\,, \quad b' \in \{\mathtt{tt}, \mathtt{ff}\} \right\}\,; \tag{17}$$

in case $b = \mathtt{tt}$,

$$\left\{ \left(\begin{Bmatrix} \big(q_1^1,\, F(q_1^1)\big)\,, \;\dots,\; \big(q_{l_1}^1,\, F(q_{l_1}^1)\big), \\ \vdots \\ \big(q_1^n,\, F(q_1^n)\big)\,, \;\dots,\; \big(q_{l_n}^n,\, F(q_{l_n}^n)\big) \end{Bmatrix}, \begin{array}{c} \min\{v, u^1, \dots, u^n\}\,, \\ b' \end{array} \right) \;\middle|\; \begin{array}{l} \big((q_1^i \wedge \dots \wedge q_{l_i}^i) \wedge u^i \big) \\ \qquad \in \delta(q^i, a)\ , \\ b' \in \{\mathtt{tt}, \mathtt{ff}\} \end{array} \right\}\,. \tag{18}$$

In each case ($b = \mathtt{ff}$ or $\mathtt{tt}$), different a-successors of $\widetilde{q}$ arise from: 1) different choices of a disjunct of a DNF formula $\delta(q^i, a)$, for $i \in [0, n]$; and 2) different choices of b' (it can always be chosen from $\mathtt{tt}$ and $\mathtt{ff}$).

In the setting of [18] (that is Boolean instead of quantitative), the state space Q' of the nondeterministic automaton obtained as a translation of an alternating one is $\mathcal{P}(Q \times \{0, 1\})$. Its quantitative adaptation $\mathcal{P}(Q \times V_Q)$ occurs as the first component of Q' in our above quantitative construction; the rest $V_\delta \times \{\mathtt{ff}, \mathtt{tt}\}$ of Q' is there for handling quantitative acceptance.

It is not hard to see that $\mathcal{A}$ and $\mathcal{A}'$ have the same language.[3] For example, in a state $\widetilde{q} = \big(\{\!\{ (q^1, v^1), \dots, (q^n, v^n) \}\!\}, v, b \big)$ of $\mathcal{A}'$:

- The pair (q^i, v^i) is that of the *current state* and what we call the *internally accumulated acceptance value*.
- The set $\{\!\{ (q^1, v^1), \dots, (q^n, v^n) \}\!\}$ stands for the *conjunction* of these pairs.
- The second component $v \in [0, 1]$ of $\widetilde{q}$ is for keeping track of: the values at the leaves of the corresponding run tree, more precisely the smallest among such.
- The flag $b \in \{\mathtt{ff}, \mathtt{tt}\}$ is called an *exposition flag*: it determines if the internally accumulated acceptance values $v^1, \dots, v^n$ should be exposed or not. Note the definition of F': the acceptance value of a state of $\mathcal{A}'$ is nonzero only if the exposition flag b is $\mathtt{tt}$.

Let us comment on the definition of the transition function δ'. Starting from $\widetilde{q} = \big(\{\!\{ (q^1, v^1), \dots, (q^n, v^n) \}\!\}, v, b \big)$—in which the "current state" is the conjunction $q^1 \wedge q^2 \wedge \dots \wedge q^n$—we choose one disjunct $q_1^i \wedge \dots \wedge q_{l_i}^i \in \delta(q^i, a)$ for each q^i and the "next state" is

$$(q_1^1 \wedge \dots \wedge q_{l_1}^1) \wedge (q_1^2 \wedge \dots \wedge q_{l_2}^2) \wedge \dots \wedge (q_1^n \wedge \dots \wedge q_{l_n}^n)\ .$$

If the exposition flag b is $\mathtt{ff}$ then we keep accumulating the acceptance values that we have seen since the last exposition, resulting in the occurrence of max in (17). If the flag is $\mathtt{tt}$ then the internally accumulated acceptance values are "used" (see the definition of F'), and

[3] A more rigorous proof can be given via formulating an acceptance game for an alternating $[0, 1]$-acceptance automaton.

these values must be "forgotten" so that we simulate a Büchi-like acceptance condition for $\mathcal{A}$. Therefore in (18), there are no $v^1, \dots, v^n$ occurring and we have a fresh start. ◀

The state space Q' of $\mathcal{A}'$ in the previous proof can actually be smaller: we can identify two states (Y, v, b) and (Y', v, b) if $\min\{v' \in V_Q \mid (q, v') \in Y\} = \min\{v' \in V_Q \mid (q, v') \in Y'\}$ holds for each $q \in Q$—this is the case for example when $Y = \{(q, \frac{1}{2}), (q, 1)\}$ and $Y' = \{(q, \frac{1}{2})\}$. Therefore we only need states (Y, v, b) such that $\forall (q, v), (q', v') \in Y.\, (q = q' \Rightarrow v = v')$, that is, Y can be regarded as a partial function. Summarizing, we can reduce the state space to $(V_Q \cup \{*\})^Q \times V_\delta \times \{\mathtt{ff}, \mathtt{tt}\}$. The size of the first component is $2^{|Q| \times \log |V_Q|}$, while it was $2^{|Q| \times |V_Q|}$ before this optimization.

B.2 Proof of Prop. 3.5

The proof is an adaptation of that of Prop. 3.4. Here we combine the usual construction of synchronous products of automata, with the idea of exposition flags.

Proof. Let $\mathcal{A}_i = (\Sigma, Q_i, I_i, \delta_i, F_i)$ for each $i \in \{1, \dots, k\}$. We define $f(\mathcal{A}_1, \dots, \mathcal{A}_k) = (\Sigma, Q, I, \delta, F)$ as follows. Its state space Q is $Q = (\prod_{1 \le i \le k} (Q_i \times V_i)) \times \{\mathtt{ff}, \mathtt{tt}\}$ where $V_i = \{0\} \cup \{v \in [0,1] \mid \exists q \in Q_i.\, F_i(q) = v\}$. The set I of initial states is $I = \{((q_1, 0), \dots, (q_k, 0), \mathtt{ff}) \mid q_1 \in I_1, \dots, q_k \in I_k\}$. The acceptance function is defined by

$$F\big((q_1, v_1), \dots, (q_k, v_k), b\big) \;=\; \begin{cases} f(v_1, \dots, v_k) & \text{if } b = \mathtt{tt} \\ 0 & \text{otherwise.} \end{cases} \tag{19}$$

The transition function $\delta \colon Q \times \Sigma \to \mathcal{P}(Q)$ is defined as follows. Let $q = \big((q_1, v_1), \dots, (q_k, v_k), b\big) \in Q$, and $a \in \Sigma$.

$$\delta(q, a) = \begin{cases} \left(\prod_{1 \le i \le k} \big\{\, (q_i', F_i(q_i')) \,\big|\, q_i' \in \delta_i(q_i, a) \,\big\} \right) \times \{\mathtt{ff}, \mathtt{tt}\} & \text{if } b = \mathtt{tt} \\ \left(\prod_{1 \le i \le k} \big\{\, (q_i', \max\{v_i, F_i(q_i')\}) \,\big|\, q_i' \in \delta_i(q_i, a) \,\big\} \right) \times \{\mathtt{ff}, \mathtt{tt}\} & \text{otherwise.} \end{cases} \tag{20}$$

We shall prove that the automaton $f(\mathcal{A}_1, \dots, \mathcal{A}_k)$ indeed satisfies the requirement. Recall that, by definition, a $[0,1]$-acceptance automaton has no dead ends. Let $w \in \Sigma^\omega$ be an infinite word.

On the one hand, it follows easily from the above definition (in particular (19)) that if $\mathcal{L}(f(\mathcal{A}_1, \dots, \mathcal{A}_k))(w) = \overline{v}$, there exist $\overline{v}_1, \dots, \overline{v}_k \in [0,1]$ such that: $f(\overline{v}_1, \dots, \overline{v}_k) = \overline{v}$, and $\mathcal{L}(\mathcal{A}_i)(w) \ge \overline{v}_i$ for each $i \in \{1, \dots, k\}$. Hence the monotonicity of f yields $\mathcal{L}\big(f(\mathcal{A}_1, \dots, \mathcal{A}_k)\big)(w) \le f\big(\mathcal{L}(\mathcal{A}_1)(w), \dots, \mathcal{L}(\mathcal{A}_k)(w)\big)$.

On the other hand, assuming that $\mathcal{L}(\mathcal{A}_i)(w) = \overline{v}_i$ for each $i \in \{1, \dots, k\}$, it is not hard to see that $\mathcal{L}\big(f(\mathcal{A}_1, \dots, \mathcal{A}_k)\big)(w) \ge f(\overline{v}_1, \dots, \overline{v}_k) = f\big(\mathcal{L}(\mathcal{A}_1)(w), \dots, \mathcal{L}(\mathcal{A}_k)(w)\big)$. Here the intuition about the automaton $f(\mathcal{A}_1, \dots, \mathcal{A}_k)$, and especially its state $q = \big((q_1, v_1), \dots, (q_k, v_k), b\big) \in Q$, is as follows.

- The automaton $f(\mathcal{A}_1, \dots, \mathcal{A}_k)$ is essentially a synchronous product of $\mathcal{A}_1, \dots, \mathcal{A}_k$; the state $q_i \in Q_i$ is the current state of the constituent automaton $\mathcal{A}_i$.
- Each constituent automaton $\mathcal{A}_i$ is additionally equipped with a register for storing "the greatest acceptance value that is recently seen." The value v_i is the one stored in that register.

- The flag $b \in \{\mathtt{ff}, \mathtt{tt}\}$ decides if the stored acceptance value v_i is "exposed" or not. See (19) where the acceptance value of the composed automaton $f(\mathcal{A}_1, \dots, \mathcal{A}_k)$ is nonzero only if $b = \mathtt{tt}$. Also observe that, in (20), the register v_i is reset to the current acceptance value $F_i(q_i')$ when the register is exposed (i.e. $b = \mathtt{tt}$).

Following this intuition, it is not hard to see that the claimed fact $\mathcal{L}\big(f(\mathcal{A}_1, \dots, \mathcal{A}_k)\big)(w) \geq f(\overline{v}_1, \dots, \overline{v}_k)$ is witnessed by a run such that: it does not expose the register values before all the registers acquire the values $\overline{v}_1, \dots, \overline{v}_k$; and once they have all done so, the register values are exposed by setting $b = \mathtt{tt}$.

From the above two inequalities, we conclude that $\mathcal{L}\big(f(\mathcal{A}_1, \dots, \mathcal{A}_k)\big)(w) = f\big(\mathcal{L}(\mathcal{A}_1)(w), \dots, \mathcal{L}(\mathcal{A}_k)(w)\big)$. ◀

B.3 Proof of Lem. 4.8

Proof. The state space $Q = xcl(\varphi) \times [0,1]^+$ of $\mathcal{A}^{\mathrm{p}}_{\varphi,\varepsilon}$ is infinite for three reasons: 1) the extended closure $xcl(\varphi)$ contains $\varphi_1 \,\mathsf{U}_{\eta^{+k}}\, \varphi_2$ for unbounded $k \in \mathbb{N}$ (see (5)); 2) discount factors occurring in $\vec{d} \in [0,1]^+$ are multiples of numbers from an infinite set $\{\eta(0), \eta(1), \dots\}$; and 3) the length of a discount sequence $\vec{d} \in [0,1]^+$ is potentially unbounded.

We can easily see that the reason 3) is not a problem for us: in the construction of $\mathcal{A}^{\mathrm{p}}_{\varphi,\varepsilon}$ (Def. 4.7), the length of a discount sequence $\vec{d}$ grows only when we encounter negation (i.e. in the definition of $\delta\big((\neg\psi, \vec{d}), \sigma\big)$). Therefore in a reachable state $(\psi, \vec{d})$ of $\mathcal{A}^{\mathrm{p}}_{\varphi,\varepsilon}$, the length of $\vec{d}$ is bounded by the number of negation operators occurring in φ.

To see that the reasons 1) and 2) are not problematic either, note that we obtain new states for these reasons only in the clause (11) of the definition of $\delta\big((\psi_1 \,\mathsf{U}_\eta\, \psi_2, \vec{d}), \sigma\big)$. This clause is applied only when $\eta(0) \cdot \prod_{i=1}^{n} d_i > \varepsilon$, a condition satisfied by only finitely many reachable states of $\mathcal{A}^{\mathrm{p}}$:

- The discount function η here is of the form $\eta = (\eta')^{+k}$, where η' occurs in the original formula φ and $k \in \mathbb{N}$. Since a discounting function η' tends to 0 (Def. 2.1), $\eta(0) = (\eta')^{+k}(0) = \eta'(k)$ tends to 0 as $k \to \infty$, too, making only finitely many k suitable.
- Each discount factor d_j in $\vec{d}$ is a multiple $\eta_1(k_1) \times \cdots \times \eta_m(k_m)$, where η_i is a discounting function occurring in φ and $k_i \in \mathbb{N}$. They must at least satisfy $\eta_i(k_i) > \varepsilon$: since η_i tends to 0, this allows only finitely many choices of k_i, for each η_i. Furthermore, the (necessary) condition that $d_j = \eta_1(k_1) \times \cdots \times \eta_m(k_m) > \varepsilon$ bounds the length m of the multiple, too.

◀

B.4 Proof of Lem. 4.9

Proof. In what follows let Q denote the state space of $\mathcal{A}_{\varphi,\varepsilon}$; δ denote its transition function; and F denote its acceptance function. For each $(\psi, \vec{d}) \in Q$, we define an alternation $[0,1]$-acceptance automaton $\mathcal{A}^{(\psi,\vec{d})}_{\varphi,\varepsilon}$ by changing the initial state to $(\psi, \vec{d})$, that is, $\mathcal{A}^{(\psi,\vec{d})}_{\varphi,\varepsilon} = (\mathcal{P}(AP), Q, \{(\psi, \vec{d})\}, \delta, F)$. Suppose that $\vec{d} = d_1 d_2 \dots d_n$. We prove the following more general statement, inductively on the construction of ψ:

$$\vec{d} \boxtimes [\![\pi, \psi]\!] - \varepsilon \;\leq\; \mathcal{L}(\mathcal{A}^{(\psi,\vec{d})}_{\varphi,\varepsilon})(\pi) \;\leq\; \vec{d} \boxtimes [\![\pi, \psi]\!] \tag{21}$$

for each $\pi \in (\mathcal{P}(AP))^\omega$.

The cases where $\psi = \mathtt{True}$, p, $\psi_1 \wedge \psi_2$, $\neg\psi'$ or $\mathsf{X}\psi'$ are straightforward. Here we only prove the case where $\psi = \neg\psi'$. By the definition of the automaton $\mathcal{A}_{\varphi,\varepsilon}$ we have $\mathcal{L}(\mathcal{A}^{(\neg\psi',\vec{d})}_{\varphi,\varepsilon})(\pi) =$

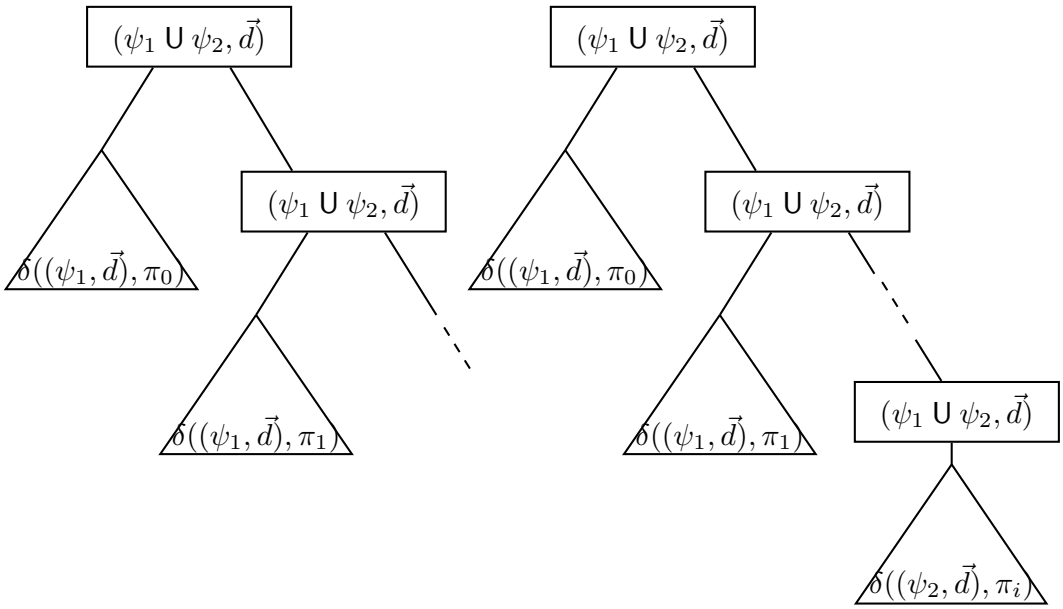

Figure 3 Possible run trees from the state $(\psi_1 \mathsf{U}\, \psi_2, \vec{d})$ in $\mathcal{A}_{\varphi,\varepsilon}$, when $|\vec{d}|$ is odd

$\mathcal{L}(\mathcal{A}_{\varphi,\varepsilon}^{(\psi', \vec{d}1)})(\pi)$, and the latter value lies in the interval $\big[\, (\vec{d}1) \boxtimes [\![\pi, \psi']\!] - \varepsilon,\ (\vec{d}1) \boxtimes [\![\pi, \psi']\!] \,\big]$ by the induction hypothesis. Now we obtain

$$(\vec{d}1) \boxtimes [\![\pi, \psi']\!] \;=\; \vec{d} \boxtimes (1 - [\![\pi, \psi']\!]) \;=\; \vec{d} \boxtimes [\![\pi, \neg\psi']\!] \ ,$$

as required. Here the former equality is due to the definition of $\boxtimes$; the latter is the semantics of $\neg\psi'$.

Suppose $\psi = \psi_1 \mathsf{U}\, \psi_2$; we first deal with the case when $|\vec{d}|$ is odd. Let $\pi \in (\mathcal{P}(AP))^\omega$. We note that, since $|\vec{d}|$ is odd, the function $\vec{d} \boxtimes (\_) \colon [0,1] \to [0,1]$ is monotone and continuous (see (7)). This is used in:

$$\begin{aligned} \vec{d} \boxtimes [\![\pi, \psi_1 \mathsf{U}\, \psi_2]\!] &= \vec{d} \boxtimes \sup_{i \in \mathbb{N}} \big\{\, \min\big\{\, [\![\pi^i, \psi_2]\!],\ \min_{0 \le j \le i-1} [\![\pi^j, \psi_1]\!] \,\big\} \,\big\} \\ &= \sup_{i \in \mathbb{N}} \big\{\, \min\big\{\, \vec{d} \boxtimes [\![\pi^i, \psi_2]\!],\ \min_{0 \le j \le i-1} \big(\vec{d} \boxtimes [\![\pi^j, \psi_1]\!]\big) \,\big\} \,\big\} \ . \end{aligned} \tag{22}$$

Now let us take a closer look at how the value $\mathcal{L}(\mathcal{A}_{\varphi,\varepsilon}^{(\psi_1 \mathsf{U} \psi_2, \vec{d})})(\pi)$ is defined for an alternating $[0,1]$-acceptance automaton $\mathcal{A}_{\varphi,\varepsilon}^{(\psi_1 \mathsf{U} \psi_2, \vec{d})}$. As seen in Def. 3.3, the notions of run tree and path are Boolean; a non-Boolean value arises for the first time as the "utility" $F^\infty(\rho)$ of a path ρ of a run tree. According to Def. 4.7 of $\mathcal{A}_{\varphi,\varepsilon}$ (in particular the definition of $\delta\big((\psi_1 \mathsf{U} \psi_2, \vec{d}), \sigma\big)$), any possible run tree τ from the state $(\psi_1 \mathsf{U}\, \psi_2, \vec{d})$ is of one of the following forms:

- the second disjunct $\delta\big((\psi_1, \vec{d}), \sigma\big) \wedge (\psi_1 \mathsf{U}\, \psi_2, \vec{d})$ is chosen all the way (Fig. 3, left), or
- the first disjunct $\delta\big((\psi_2, \vec{d}), \sigma\big)$ is eventually hit (Fig. 3, right).

In the former case, the utility $\min_{\rho \in \mathrm{path}(\tau)} F^\infty(\rho)$ of such a run tree τ is given by $\min\big\{\, F(\psi_1 \mathsf{U}\, \psi_2, \vec{d}),\ \inf_{j \in \mathbb{N}} \mathcal{L}(\mathcal{A}_{\varphi,\varepsilon}^{(\psi_1, \vec{d})})(\pi^j) \,\big\}$, where the first value $F(\psi_1 \mathsf{U} \psi_2, \vec{d})$ is induced by the rightmost path in Fig. 3, left. We have $F(\psi_1 \mathsf{U}\, \psi_2, \vec{d}) = 0$ by definition (see (12)); therefore the utility obtained in this case is 0.

In the latter case, assume that the second disjunct $\delta\big((\psi_2, \vec{d}), \sigma\big)$ is hit at depth i. The tree's utility is then given by $\min\big\{\, \mathcal{L}(\mathcal{A}_{\varphi,\varepsilon}^{(\psi_2, \vec{d})})(\pi^i),\ \min_{0 \le j \le i-1} \mathcal{L}(\mathcal{A}_{\varphi,\varepsilon}^{(\psi_1, \vec{d})})(\pi^j) \,\big\}$ where, again, the first value $\mathcal{L}(\mathcal{A}_{\varphi,\varepsilon}^{(\psi_2, \vec{d})})(\pi^i)$ arises from the rightmost path in Fig. 3, right.

Putting all these together, we have

$$\begin{aligned}
&\mathcal{L}(\mathcal{A}_{\varphi,\varepsilon}^{(\psi_1 \mathsf{U} \psi_2, \vec{d})})(\pi) \\
&= \sup_{i\in\mathbb{N}}\big(\min\big\{ \mathcal{L}(\mathcal{A}_{\varphi,\varepsilon}^{(\psi_2,\vec{d})})(\pi^i),\ \min_{0\le j\le i-1} \mathcal{L}(\mathcal{A}_{\varphi,\varepsilon}^{(\psi_1,\vec{d})})(\pi^j) \big\} \big) \\
&\in \left[\begin{array}{r} \sup_{i\in\mathbb{N}}\big(\min\big\{ \vec{d}\boxtimes[\![\pi^i,\psi_2]\!] - \varepsilon,\ \min_{0\le j\le i-1} \vec{d}\boxtimes[\![\pi^j,\psi_1]\!] - \varepsilon \big\} \big)\,, \\ \sup_{i\in\mathbb{N}}\big(\min\big\{ \vec{d}\boxtimes[\![\pi^i,\psi_2]\!],\ \min_{0\le j\le i-1} \vec{d}\boxtimes[\![\pi^j,\psi_1]\!] \big\} \big) \end{array} \right] \\
&\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad \text{by the induction hypothesis} \\
&= \big[\, \vec{d}\boxtimes[\![\pi,\psi_1 \mathsf{U} \psi_2]\!] - \varepsilon\,,\ \vec{d}\boxtimes[\![\pi,\psi_1 \mathsf{U} \psi_2]\!] \,\big] \qquad \text{by (22)},
\end{aligned}$$

as required.

Suppose that $\psi = \psi_1 \mathsf{U} \psi_2$ and that $|\vec{d}|$ is even. Let $\pi \in (\mathcal{P}(AP))^\omega$. Since $\vec{d}\boxtimes(\_)$ is antitone and continuous, the second equality below holds.

$$\begin{aligned}
\vec{d}\boxtimes[\![\pi,\psi_1 \mathsf{U} \psi_2]\!] &= \vec{d}\boxtimes \sup_{i\in\mathbb{N}}\big\{ \min\big\{ [\![\pi^i,\psi_2]\!],\ \min_{0\le j\le i-1} [\![\pi^j,\psi_1]\!] \big\} \big\} \quad \text{by def. of } [\![\pi,\psi_1 \mathsf{U} \psi_2]\!] \\
&= \inf_{i\in\mathbb{N}}\big\{ \max\big\{ \vec{d}\boxtimes[\![\pi^i,\psi_2]\!],\ \max_{0\le j\le i-1} (\vec{d}\boxtimes[\![\pi^j,\psi_1]\!]) \big\} \big\}\ .
\end{aligned} \tag{23}$$

We use the following observation. It is a quantitative adaptation of the classic duality between the temporal operators U and R ("release").

▶ **Sublemma B.1.** *Let $a_0, a_1, \dots$ and $b_0, b_1, \dots$ all be real numbers in $[0,1]$. We have*

$$\inf_{i\in\mathbb{N}}\big(\max\{b_i, \max_{0\le j\le i-1} a_j\} \big) \ = \ \Big\{ \sup_{j\in\mathbb{N}}\big(\min\{a_j, \min_{0\le i\le j} b_i\} \big), \inf_{i\in\mathbb{N}} b_i \Big\}\ ,$$

that is, denoting binary min *and* max *by $\wedge$ and $\vee$:*

$$\inf_{i\in\mathbb{N}}\big(b_i \vee (a_0 \vee a_1 \vee \cdots \vee a_{i-1}) \big) \ = \ \Big(\sup_{j\in\mathbb{N}}\big(a_j \wedge (b_0 \wedge b_1 \wedge \cdots \wedge b_j) \big) \Big) \vee \inf_{i\in\mathbb{N}} b_i\ . \tag{24}$$

Proof. (Of Sublem. B.1) We distinguish two cases. Let us first assume that there exists $i \in \mathbb{N}$ such that $b_i < a_0 \vee a_1 \vee \cdots \vee a_{i-1}$. Let k be the least number among such, that is, k satisfies that

$$b_k < a_0 \vee a_1 \vee \cdots \vee a_{k-1} \quad \text{and} \quad \forall i \in [0, k-1].\ b_i \ge a_0 \vee a_1 \vee \cdots \vee a_{i-1}\ . \tag{25}$$

Moreover, let $l \in [0, k-1]$ be a number such that $a_l = a_0 \vee a_1 \vee \cdots \vee a_{k-1}$. We have

$$\begin{aligned}
&\inf_{i\in\mathbb{N}}\big(b_i \vee (a_0 \vee a_1 \vee \cdots \vee a_{i-1}) \big) \\
&= b_0 \ \wedge\ (b_1 \vee a_0) \ \wedge\ (b_2 \vee a_0 \vee a_1) \ \wedge\ \cdots \\
&\qquad\qquad \wedge\ (b_k \vee a_0 \vee \cdots \vee a_{k-1}) \wedge \inf_{i\ge k+1}\big(b_i \vee (a_0 \vee a_1 \vee \cdots \vee a_{i-1}) \big) \\
&= b_0 \wedge b_1 \wedge \cdots \wedge b_{k-1} \wedge a_l \ \wedge\ \inf_{i\ge k+1}\big(b_i \vee (a_0 \vee a_1 \vee \cdots \vee a_{i-1}) \big) \quad \text{by def. of } k,\ (25).
\end{aligned}$$

Since we have $a_l \le b_i \vee (a_0 \vee a_1 \vee \cdots \vee a_{i-1})$ for each $i \in [k+1, \infty)$,

$$a_l \ \le\ \inf_{i\ge k+1}\big(b_i \vee (a_0 \vee a_1 \vee \cdots \vee a_{i-1}) \big)$$

and we obtain

$$\inf_{i\in\mathbb{N}}\big(b_i \vee (a_0 \vee a_1 \vee \cdots \vee a_{i-1}) \big) \ = \ b_0 \wedge b_1 \wedge \cdots \wedge b_{k-1} \wedge a_l\ . \tag{26}$$

Now we compare the last value $b_0 \wedge b_1 \wedge \cdots \wedge b_{k-1} \wedge a_l$ with the right-hand side of our goal (24). By the definition of k and l, for each $j \in [0, k-1]$, we have

$$a_j \leq a_0 \vee a_1 \vee \cdots \vee a_{k-1} = a_l \quad \text{and} \quad \forall i \in [j+1, k-1].\ a_j \leq a_0 \vee a_1 \vee \cdots \vee a_{i-1} \leq b_i \ ,$$

yielding

$$\begin{aligned} a_j &\leq a_l \wedge b_{j+1} \wedge b_{j+2} \wedge \cdots \wedge b_{k-1} \ , \quad \text{and hence} \\ a_j \wedge (b_0 \wedge b_1 \wedge \cdots \wedge b_j) &\leq b_0 \wedge b_1 \wedge \cdots \wedge b_j \wedge b_{j+1} \wedge \cdots \wedge b_{k-1} \wedge a_l \ . \end{aligned}$$

The last inequality holds for each $j \in [k, \infty)$, too:

$$\begin{aligned} a_j \wedge (b_0 \wedge b_1 \wedge \cdots \wedge b_j) &\leq b_0 \wedge b_1 \wedge \cdots \wedge b_{k-1} \wedge b_k \\ &\leq b_0 \wedge b_1 \wedge \cdots \wedge b_{k-1} \wedge (a_0 \vee a_1 \vee \cdots \vee a_{k-1}) \\ &\qquad \text{by def. of } k, (25) \\ &= b_0 \wedge b_1 \wedge \cdots \wedge b_{k-1} \wedge a_l \quad \text{by def. of } l. \end{aligned}$$

Consequently

$$\sup_{j \in \mathbb{N}} \big(a_j \wedge (b_0 \wedge b_1 \wedge \cdots \wedge b_j) \big) \leq b_0 \wedge b_1 \wedge \cdots \wedge b_{k-1} \wedge a_l \ . \tag{27}$$

We turn to the other part $\inf_{i \in \mathbb{N}} b_i$ of the right-hand side of (24). By the definition of k and l, we have $b_k \leq a_0 \vee a_1 \vee \cdots \vee a_{k-1} = a_l$. Therefore

$$\inf_{i \in \mathbb{N}} b_i \leq b_0 \wedge b_1 \wedge \cdots \wedge b_{k-1} \wedge b_k \leq b_0 \wedge b_1 \wedge \cdots \wedge b_{k-1} \wedge a_l \ . \tag{28}$$

By (27) and (28),

$$\Big(\sup_{j \in \mathbb{N}} \big(a_j \wedge (b_0 \wedge b_1 \wedge \cdots \wedge b_j) \big) \Big) \vee \inf_{i \in \mathbb{N}} b_i \leq b_0 \wedge b_1 \wedge \cdots \wedge b_{k-1} \wedge a_l \ , \tag{29}$$

on the one hand. On the other hand, since $l \in [0, k-1]$,

$$\begin{aligned} \Big(\sup_{j \in \mathbb{N}} \big(a_j \wedge (b_0 \wedge b_1 \wedge \cdots \wedge b_j) \big) \Big) \vee \inf_{i \in \mathbb{N}} b_i &\geq \sup_{j \in \mathbb{N}} \big(a_j \wedge (b_0 \wedge b_1 \wedge \cdots \wedge b_j) \big) \\ &\geq a_l \wedge (b_0 \wedge b_1 \wedge \cdots \wedge b_l) \\ &\geq a_l \wedge (b_0 \wedge b_1 \wedge \cdots \wedge b_{k-1}) \ . \end{aligned} \tag{30}$$

By (29) and (30),

$$\Big(\sup_{j \in \mathbb{N}} \big(a_j \wedge (b_0 \wedge b_1 \wedge \cdots \wedge b_j) \big) \Big) \vee \inf_{i \in \mathbb{N}} b_i = b_0 \wedge b_1 \wedge \cdots \wedge b_{k-1} \wedge a_l \ . \tag{31}$$

By (26) and (31),

$$\inf_{i \in \mathbb{N}} \big(b_i \vee (a_0 \vee a_1 \vee \cdots \vee a_{i-1}) \big) = \Big(\sup_{j \in \mathbb{N}} \big(a_j \wedge (b_0 \wedge b_1 \wedge \cdots \wedge b_j) \big) \Big) \vee \inf_{i \in \mathbb{N}} b_i \ .$$

This establish the claim, in our first case where there exists $i \in \mathbb{N}$ such that $b_i < a_0 \vee a_1 \vee \cdots \vee a_{i-1}$.

In the other case we assume that $b_i \geq a_0 \vee a_1 \vee \cdots \vee a_{i-1}$ for each $i \in \mathbb{N}$. By this assumption, $b_i \vee (a_0 \vee a_1 \vee \cdots \vee a_{i-1}) = b_i$ for each $i \in \mathbb{N}$. Therefore

$$\inf_{i \in \mathbb{N}} \big(b_i \vee (a_0 \vee a_1 \vee \cdots \vee a_{i-1}) \big) = \inf_{i \in \mathbb{N}} b_i \ . \tag{32}$$

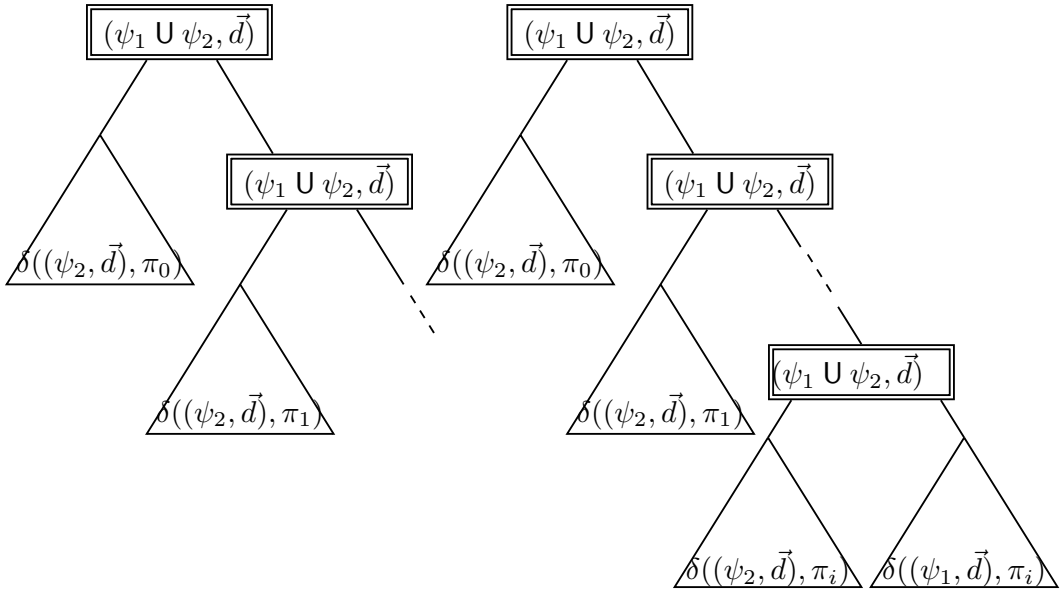

Figure 4 Possible run trees from the state $(\psi_1 \mathsf{U} \psi_2, \vec{d})$ in $\mathcal{A}_{\varphi,\varepsilon}$, when $|\vec{d}|$ is even. The double-lined nodes have the acceptance value 1.

Let us now fix $j \in \mathbb{N}$. For each $i \in [j+1, \infty)$ we have $a_j \le a_0 \vee a_1 \vee \cdots \vee a_{i-1} \le b_i$, where the latter inequality holds because of the assumption. Therefore $a_j \le \inf_{i \ge j+1} b_i$; this is used in

$$a_j \wedge (b_0 \wedge b_1 \wedge \cdots \wedge b_j) \;\le\; (\inf_{i \ge j+1} b_i) \wedge (b_0 \wedge b_1 \wedge \cdots \wedge b_j) \;=\; \inf_{i \in \mathbb{N}} b_i \;.$$

This holds for any $j \in \mathbb{N}$; therefore $\sup_{j \in \mathbb{N}} \big(a_j \wedge (b_0 \wedge b_1 \wedge \cdots \wedge b_j) \big) \le \inf_{i \in \mathbb{N}} b_i$. This yields $\Big(\sup_{j \in \mathbb{N}} \big(a_j \wedge (b_0 \wedge b_1 \wedge \cdots \wedge b_j) \big) \Big) \vee \inf_{i \in \mathbb{N}} b_i = \inf_{i \in \mathbb{N}} b_i$, which is combined with (32) and proves the claim (24). This concludes the proof of Sublem. B.1. ◀

We turn back to the proof of Lem. 4.9. By letting $a_j = \vec{d} \boxtimes [\![\pi^j, \psi_1]\!]$ and $b_i = \vec{d} \boxtimes [\![\pi^i, \psi_2]\!]$ in Sublem. B.1, we obtain

$$\begin{aligned} &\inf_{i \in \mathbb{N}} \big(\max\{ \vec{d} \boxtimes [\![\pi^i, \psi_2]\!], \max_{0 \le j \le i-1} (\vec{d} \boxtimes [\![\pi^j, \psi_1]\!]) \} \big) \\ &= \max\{ \sup_{j \in \mathbb{N}} \big(\min\{ \vec{d} \boxtimes [\![\pi^j, \psi_1]\!], \min_{0 \le i \le j} (\vec{d} \boxtimes [\![\pi^i, \psi_2]\!]) \} \big), \inf_{i \in \mathbb{N}} (\vec{d} \boxtimes [\![\pi^i, \psi_2]\!]) \} \;. \end{aligned} \tag{33}$$

By (23) and (33), we have

$$\begin{aligned} &\vec{d} \boxtimes [\![\pi, \psi_1 \mathsf{U} \psi_2]\!] \\ &= \max\{ \sup_{j \in \mathbb{N}} \{ \min\{ \vec{d} \boxtimes [\![\pi^j, \psi_1]\!], \min_{0 \le i \le j} (\vec{d} \boxtimes [\![\pi^i, \psi_2]\!]) \} \}, \inf_{i \in N} (\vec{d} \boxtimes [\![\pi^i, \psi_2]\!]) \} \;. \end{aligned} \tag{34}$$

Let us now look at the value $\mathcal{L}(\mathcal{A}_{\varphi,\varepsilon}^{(\psi_1 \mathsf{U} \psi_2, \vec{d})})(\pi)$. We analyze possible run trees τ starting from the state $(\psi_1 \mathsf{U} \psi_2, \vec{d})$, much like in the previous case where $|\vec{d}|$ is odd (in the current case it is even). It is easily seen from Def. 4.7 that τ is of one of the forms shown in Fig. 4.

- If τ is of the form in Fig. 4 on the left, its utility $\min_{\rho \in \mathrm{path}(\tau)} F^\infty(\rho)$ is $\inf_{j \in \mathbb{N}} \mathcal{L}(\mathcal{A}_{\varphi,\varepsilon}^{(\psi_2, \vec{d})})(\pi^j)$; note that the rightmost path's value of F^∞ is 1 and hence does not appear here.

- If τ is of the form in Fig. 4 on the right, its utility $\min_{\rho\in\mathrm{path}(\tau)} F^{\infty}(\rho)$ is given by $\min\{\, \mathcal{L}(\mathcal{A}^{(\psi_1,\vec{d})}_{\varphi,\varepsilon})(\pi^i),\ \min_{0\le j\le i} \mathcal{L}(\mathcal{A}^{(\psi_2,\vec{d})}_{\varphi,\varepsilon})(\pi^j) \,\}$ where i is the depth of the last occurrence of the node $(\psi_1 \mathsf{U} \psi_2, \vec{d})$.

The value $\mathcal{L}(\mathcal{A}^{(\psi_1 \mathsf{U}\psi_2,\vec{d})}_{\varphi,\varepsilon})(\pi)$ is defined as the supremum of these utilities. Therefore:

$$
\begin{aligned}
&\mathcal{L}(\mathcal{A}^{(\psi_1 \mathsf{U}\psi_2,\vec{d})}_{\varphi,\varepsilon})(\pi)\\
&= \max\Big\{ \sup_{i\in\mathbb{N}} \Big(\min\big\{ \mathcal{L}(\mathcal{A}^{(\psi_1,\vec{d})}_{\varphi,\varepsilon})(\pi^i),\ \max_{0\le j\le i} \mathcal{L}(\mathcal{A}^{(\psi_2,\vec{d})}_{\varphi,\varepsilon})(\pi^j) \big\} \Big),\ \inf_{j\in\mathbb{N}} \mathcal{L}(\mathcal{A}^{(\psi_2,\vec{d})}_{\varphi,\varepsilon})(\pi^j) \Big\}\\
&\in \left[\begin{array}{l} \max\Big\{ \sup_{j\in\mathbb{N}} \Big(\min\big\{ \vec{d}\boxtimes[\![\pi^j,\psi_1]\!],\ \min_{0\le i\le j} \vec{d}\boxtimes[\![\pi^i,\psi_2]\!] \big\} \Big) ,\ \inf_{i\in\mathbb{N}} \vec{d}\boxtimes[\![\pi^i,\psi_2]\!] \Big\} - \varepsilon\,, \\ \qquad \max\Big\{ \sup_{j\in\mathbb{N}} \Big(\min\big\{ \vec{d}\boxtimes[\![\pi^j,\psi_1]\!],\ \min_{0\le i\le j} \vec{d}\boxtimes[\![\pi^i,\psi_2]\!] \big\} \Big) ,\ \inf_{i\in\mathbb{N}} \vec{d}\boxtimes[\![\pi^i,\psi_2]\!] \Big\} \end{array} \right]\\
&\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad \text{by the induction hypothesis}\\
&= \big[\, \vec{d}\boxtimes[\![\pi,\psi_1 \mathsf{U} \psi_2]\!] - \varepsilon\,,\ \vec{d}\boxtimes[\![\pi,\psi_1 \mathsf{U} \psi_2]\!] \,\big] \qquad \text{by (34)},
\end{aligned}
$$

concluding the case when $\psi = \psi_1 \mathsf{U} \psi_2$ and $|\vec{d}|$ is even.

Suppose that $\psi = \psi_1 \mathsf{U}_{\eta^{+k}} \psi_2$ and that $|\vec{d}|$ is odd. We prove the claim by induction on k, going backwards, decrementing k starting from the event horizon towards $k = 0$. As the base case, assume that k is big enough and we are beyond the event horizon, that is, $\eta(k) \cdot \prod_{i=1}^{n} d_i \le \varepsilon$. Let $\pi \in (\mathcal{P}(AP))^{\omega}$ and $\vec{d} = d_1 \dots d_n$. Then we have $[\![\pi,\psi]\!] \cdot \prod_{i=0}^{n} d_i = [\![\pi,\psi_1 \mathsf{U}_{\eta^{+k}} \psi_2]\!] \cdot \prod_{i=0}^{n} d_i \le \varepsilon$, by Lem. 2.6 and that $\eta^{+k}(0) = \eta(k)$. It follows from (7) that we have $0 \le \vec{d}\boxtimes[\![\pi,\psi]\!] - \vec{d}\boxtimes 0 \le \varepsilon$ (note that $n = |\vec{d}|$ is odd). Therefore

$$
\begin{aligned}
\mathcal{L}(\mathcal{A}^{(\psi_1 \mathsf{U}_{\eta^{+k}}\psi_2,\vec{d})}_{\varphi,\varepsilon})(\pi) &= \vec{d}\boxtimes 0 \qquad \text{by (10)}\\
&\in \big[\, \vec{d}\boxtimes[\![\pi,\psi]\!] - \varepsilon,\ \vec{d}\boxtimes[\![\pi,\psi]\!] \,\big]\ .
\end{aligned}
$$

Now, as the step case, assume that $\eta(k) \cdot \prod_{i=1}^{n} d_i > \varepsilon$ and that the claim has been shown for $k+1$. The analogue below of $\psi_1 \mathsf{U} \psi_2 \cong \psi_2 \vee \big(\psi_1 \wedge \mathsf{X}(\psi_1 \mathsf{U} \psi_2)\big)$ follows easily from Def. 2.5:

$$
[\![\pi,\psi_1 \mathsf{U}_{\eta^{+k}} \psi_2]\!] \;=\; \max\big\{\, \eta^{+k}(0)\cdot[\![\pi,\psi_2]\!],\ \min\big\{\, \eta^{+k}(0)\cdot[\![\pi,\psi_1]\!],\ [\![\pi^1,\psi_1 \mathsf{U}_{\eta^{+(k+1)}} \psi_2]\!] \,\big\} \,\big\}\ .
$$

Therefore

$$
\begin{aligned}
&\vec{d}\boxtimes[\![\pi,\psi_1 \mathsf{U}_{\eta^{+k}} \psi_2]\!]\\
&= \max\big\{\, \vec{d}\boxtimes\big(\eta^{+k}(0)\cdot[\![\pi,\psi_2]\!]\big),\\
&\qquad \min\big\{\, \vec{d}\boxtimes\big(\eta^{+k}(0)\cdot[\![\pi,\psi_1]\!]\big),\ \vec{d}\boxtimes[\![\pi^1,\psi_1 \mathsf{U}_{\eta^{+(k+1)}} \psi_2]\!] \,\big\} \,\big\} \qquad (35)\\
&= \max\big\{\, \big(\vec{d}\odot\eta^{+k}(0)\big)\boxtimes[\![\pi,\psi_2]\!],\\
&\qquad \min\big\{\, \big(\vec{d}\odot\eta^{+k}(0)\big)\boxtimes[\![\pi,\psi_1]\!],\ \vec{d}\boxtimes[\![\pi^1,\psi_1 \mathsf{U}_{\eta^{+(k+1)}} \psi_2]\!] \,\big\} \,\big\}\ ,
\end{aligned}
$$

where the first equality is due to the monotonicity of $\vec{d}\boxtimes(\_)$, and the second is by (8). Now

$$
\begin{aligned}
&\mathcal{L}(\mathcal{A}^{(\psi_1 \mathsf{U}_{\eta^{+k}}\psi_2,\vec{d})}_{\varphi,\varepsilon})(\pi)\\
&= \max\big\{\, \mathcal{L}(\mathcal{A}^{(\psi_2,\vec{d}\odot\eta^{+k}(0))}_{\varphi,\varepsilon})(\pi),\ \min\big\{\, \mathcal{L}(\mathcal{A}^{(\psi_1,\vec{d}\odot\eta^{+k}(0))}_{\varphi,\varepsilon})(\pi),\ \mathcal{L}(\mathcal{A}^{(\psi_1 \mathsf{U}_{\eta^{+(k+1)}}\psi_2,\vec{d})}_{\varphi,\varepsilon})(\pi^1) \,\big\} \,\big\}
\end{aligned}
$$

by Def. 4.7. By the induction hypothesis (the claim has been shown for simpler formulas as

well as $\psi_1 \mathsf{U}_{\eta^{+(k+1)}} \psi_2$), a lower bound of the above value is given by

$$\begin{aligned}
&\max\Big\{ \Big((\vec{d} \odot \eta^{+k}(0)) \boxtimes [\![\pi, \psi_2]\!] \Big) - \varepsilon, \\
&\qquad \min\Big\{ \Big((\vec{d} \odot \eta^{+k}(0)) \boxtimes [\![\pi, \psi_1]\!] \Big) - \varepsilon,\ \Big(\vec{d} \boxtimes [\![\pi^1, \psi_1 \mathsf{U}_{\eta^{+(k+1)}} \psi_2]\!] \Big) - \varepsilon \Big\} \Big\} \\
&= \vec{d} \boxtimes [\![\pi, \psi_1 \mathsf{U}_{\eta^{+k}} \psi_2]\!] - \varepsilon \qquad \text{by (35).}
\end{aligned}$$

Similarly an upper bound $\vec{d} \boxtimes [\![\pi, \psi_1 \mathsf{U}_{\eta^{+k}} \psi_2]\!]$ is obtained by the induction hypothesis and (35). This proves the claim.

The remaining case where $\psi = \psi_1 \mathsf{U}_{\eta^{+k}} \psi_2$ and $|\vec{d}|$ is even is similar to the last case. We describe only the base case of induction, where k is big enough so that $\eta(k) \cdot \prod_{i=1}^{n} d_i \leq \varepsilon$. By Lem. 2.6 we have $[\![\pi, \psi]\!] \in [0, \eta^k(0)]$; therefore

$$0 \ \leq\ \eta(k) - [\![\pi, \psi]\!] \ \leq\ \eta(k) \ \leq\ \varepsilon / \textstyle\prod_{i=1}^{n} d_i \ .$$

By (7) and that n is even, we have

$$\vec{d} \boxtimes [\![\pi, \psi]\!] - \vec{d} \boxtimes \eta(k) \ =\ \big(\prod_{i=1}^{n} d_i\big) \cdot \big(\eta(k) - [\![\pi, \psi]\!] \big) \ \in\ [0, \varepsilon] \ .$$

Hence

$$\begin{aligned}
\mathcal{L}(\mathcal{A}_{\varphi,\varepsilon}^{(\psi_1 \mathsf{U}_{\eta^{+k}} \psi_2, \vec{d})})(\pi) &= \vec{d} \boxtimes \eta(k) \qquad \text{by (10)} \\
&\in \big[\, \vec{d} \boxtimes [\![\pi, \psi]\!] - \varepsilon,\ \vec{d} \boxtimes [\![\pi, \psi]\!] \,\big] \ .
\end{aligned}$$

This concludes the proof. ◀

B.5 Proof of Lem. 4.12

Proof. It follows easily from the definition that there is a bijective correspondence between: a run $\zeta = (q'_0, s'_0) \bullet (q'_1, s'_1) \bullet \dots$ of $\mathcal{A} \times \mathcal{K}$; and a pair (ξ, ρ) of a path $\xi = s'_0 s'_1 \dots \in \mathrm{path}(\mathcal{K})$ of $\mathcal{K}$ and a run ρ over $\lambda(\xi)$ of $\mathcal{A}$. Moreover, the acceptance value of ζ in $\mathcal{A} \times \mathcal{K}$ is equal to that of ρ in $\mathcal{A}$. The claim follows immediately. ◀

B.6 Proof of Thm. 4.13

Proof.

$$\begin{aligned}
[\![s_0 s_1 \dots, \varphi]\!] \ &\geq\ \mathcal{L}(\mathcal{A}^{\mathrm{na}}_{\varphi,\varepsilon})\big(\, \lambda(s_0)\lambda(s_1)\dots \big) \qquad \text{by Cor. 4.10} \\
&=\ \max_{\xi \in \mathrm{path}(\mathcal{K})} \mathcal{L}(\mathcal{A}^{\mathrm{na}}_{\varphi,\varepsilon})\big(\, \lambda(\xi) \,\big) \qquad \text{by Lem. 4.12} \\
&\geq\ \sup_{\xi \in \mathrm{path}(\mathcal{K})} [\![\xi, \varphi]\!] - \varepsilon \qquad \text{by Cor. 4.10.}
\end{aligned}$$

The solution $s_0 s_1 \dots$ thus obtained arises from a lasso computation of $\mathcal{A}^{\mathrm{na}}_{\varphi,\varepsilon} \times \mathcal{K}$ (by the algorithm in Lem. 3.2), hence is ultimately periodic. ◀

B.7 Proof of Prop. 4.15

Proof. In the proof of Lem. A.1, we construct $\mathcal{A}^{\vec{d}}_{\varphi,\varepsilon}$ inductively. We shall therefore prove, inductively on the construction on φ, that the size of the state space of $\mathcal{A}^{\vec{d}}_{\varphi,\varepsilon}$ is singly exponential in $|\langle\varphi\rangle|$ and in the length of the description of ε.

In the case where $\varphi = \mathtt{True}, p, \varphi_1 \wedge \varphi_2, \neg\varphi', \mathsf{X}\varphi'$ or $\varphi_1 \,\mathsf{U}\, \varphi_2$, the claim is obvious.

Suppose that $\varphi = \varphi_1 \,\mathsf{U}_{\exp_\lambda^{+k}}\, \varphi_2$ where $\lambda \in (0,1)$. Let $k_{\max} = \lceil \log_\lambda \varepsilon \rceil + 1$. Recall that the construction of $\mathcal{A}^{\vec{d}}_{\varphi_1 \mathsf{U}_{\exp_\lambda^{+k}} \varphi_2, \varepsilon}$ in Lem. A.1 is by backward induction on k, from $k = k_{\max}$ to $k = 0$. In the base case when $k = k_{\max}$, we have $\exp_\lambda^{+k}(0) \leq \varepsilon$ (beyond the event horizon); in this case the size of the state space of $\mathcal{A}^{\vec{d}}_{\varphi,\varepsilon}$ is one. In the step case, the state space of $\mathcal{A}^{\vec{d}}_{\varphi_1 \mathsf{U}_{\exp_\lambda^{+k}} \varphi_2, \varepsilon}$ is the union of: those of the two automata for φ_1 and φ_2; that of the automaton $\mathcal{A}^{\vec{d}}_{\varphi_1 \mathsf{U}_{\exp_\lambda^{+(k+1)}} \varphi_2, \varepsilon}$; and the singleton of the initial state of $\mathcal{A}^{\vec{d}}_{\varphi,\varepsilon}$. Overall, the state space of $\mathcal{A}^{\vec{d}}_{\varphi_1 \mathsf{U}_{\exp_\lambda^{+k}} \varphi_2, \varepsilon}$ increases as k decreases, and the maximum is when $k = 0$—in which case the state space of $\mathcal{A}^{\vec{d}}_{\varphi_1 \mathsf{U}_{\exp_\lambda} \varphi_2, \varepsilon}$ is roughly $\mathcal{O}(k_{\max}) = \mathcal{O}(\lceil \log_\lambda \varepsilon \rceil + 1)$ copies of those of the two automata for φ_1 and φ_2. Now we appeal to the fact used in [3] that the value $k_{\max} \sim \log_\lambda \varepsilon = \log \varepsilon / \log \lambda$ is polynomial in the length of the description of λ—hence in $|\langle\varphi\rangle|$—and ε.[4] By this fact and the induction hypothesis, the size of the state space of $\mathcal{A}^{\vec{d}}_{\varphi,\varepsilon}$ is singly exponential in $|\langle\varphi\rangle|$ and in the length of the description of ε.

Suppose that $\varphi = \varphi_1 \oplus \varphi_2$. Since $(\vec{d} \boxtimes v_1 - \varepsilon) \oplus (\vec{d} \boxtimes v_2 - \varepsilon) = \vec{d} \boxtimes (v_1 \oplus v_2) - \varepsilon$, we have $\mathcal{A}^{\vec{d}}_{\varphi,\varepsilon}$ coincide with $\mathcal{A}^{\vec{d}}_{\varphi_1,\varepsilon} \oplus \mathcal{A}^{\vec{d}}_{\varphi_2,\varepsilon}$—where the latter is defined in Prop. 3.5. (We note that the construction in Prop. 3.5 can be readily adapted to *alternating* $[0,1]$-acceptance automata, too.) Hence the size of the state space of $\mathcal{A}^{\vec{d}}_{\varphi,\varepsilon}$ is polynomial in those of $\mathcal{A}^{\vec{d}}_{\varphi_1,\varepsilon}$ and $\mathcal{A}^{\vec{d}}_{\varphi_2,\varepsilon}$. By the induction hypothesis, the size of the state space of $\mathcal{A}^{\vec{d}}_{\varphi,\varepsilon}$ is singly exponential in $|\langle\varphi\rangle|$ and in the length of the description of ε. ◀

B.8 Proof of Thm. 4.16

Proof. The construction in Prop. 3.4 (from $\mathcal{A}_{\varphi,\varepsilon}$ to $\mathcal{A}^{\mathrm{na}}_{\varphi,\varepsilon}$) results in $\mathcal{A}^{\mathrm{na}}_{\varphi,\varepsilon}$ that is exponentially bigger than $\mathcal{A}_{\varphi,\varepsilon}$; the size of the product $\mathcal{A}^{\mathrm{na}}_{\varphi,\varepsilon} \times \mathcal{K}$ (Def. 4.11) is linear in those of $\mathcal{A}^{\mathrm{na}}_{\varphi,\varepsilon}$ and $\mathcal{K}$; and finding an optimal run by Lem. 3.2 is in NLOGSPACE. Combined with Prop. 4.15, the overall complexity is EXPSPACE in $|\langle\varphi\rangle|$ and NLOGSPACE in the size of $\mathcal{K}$. ◀

B.9 Proof of Thm. 4.17

Firstly we give an alternative proof to the following statement (that is a restriction of Prop. 4.15). It is used in the proof of Thm. 4.17.

▶ **Sublemma B.2** (size of $\mathcal{A}_{\varphi,\varepsilon}$, for $\mathrm{LTL}^{\mathrm{disc}}[\mathcal{D}_{\exp}, \emptyset]$)**.** *Let φ be an $LTL^{\mathrm{disc}}[\mathcal{D}_{\exp}, \emptyset]$ formula and $\varepsilon \in (0,1) \cap \mathbb{Q}$ be a positive rational number. The size of the state space of the alternating $[0,1]$-acceptance automaton $\mathcal{A}_{\varphi,\varepsilon}$ is singly exponential in $|\langle\varphi\rangle|$ and in the length of the description of ε.* ◀

Proof. (Of Sublem. B.2) Recall that a state of $\mathcal{A}_{\varphi,\varepsilon}$ is a pair $(\psi, \vec{d})$ of $\psi \in xcl(\varphi)$ and $\vec{d} \in [0,1]^+$. We first claim that the number of different ψ's is polynomial in $|\langle\varphi\rangle|$ and $\log \varepsilon$. The claim is obvious except for the number of the formulas ψ of the form $\psi_1 \,\mathsf{U}_{\eta^{+i}}\, \psi_2$, for varying $i \in \mathbb{N}$. Let λ_0 be the maximum number in φ used as the base of an exponential

[4] It is not explicit in [3] what is meant by the description length of $\lambda \in (0,1)$. For the claimed fact to be true—that $\log_\lambda \varepsilon = \log \varepsilon / \log \lambda$ is polynomial in the length of the description of λ—we expect it to be $a + b$ where $\lambda = a/b$. For example, when $\lambda = 1 - \frac{1}{b}$, we have $\log_\lambda \varepsilon = \frac{\log \varepsilon}{\log \lambda} = \frac{\log \varepsilon}{\log(1-\frac{1}{b})} = \frac{-\log \epsilon}{\log b - \log(b-1)} \leq b \cdot (-\log \varepsilon)$ where for the last inequality we used $(\log x)' = \frac{1}{x}$. This is linear in b.

discounting function. For each subformula $\psi_1 \mathsf{U}_\eta \psi_2$ of φ, the numbers i for which we have a state $\big(\psi_1 \mathsf{U}_{\eta^{+i}} \psi_2, \vec{d}\big)$ in $\mathcal{A}_{\varphi,\varepsilon}$ is bounded by $1 + \lceil \log_{\lambda_0} \varepsilon \rceil$. Now we appeal to the fact used in [3] that the value $\log_{\lambda_0} \varepsilon = \log \varepsilon / \log \lambda_0$ is polynomial in the length of the description of λ_0—hence in $|\langle\varphi\rangle|$—and ε.

Our second claim is that the number of different $\vec{d}$'s occurring in states of $\mathcal{A}_{\varphi,\varepsilon}$ is exponential in $|\langle\varphi\rangle|$ and the description length of ε, hence is the bottleneck in complexity. The length of a discount sequence $\vec{d}$ is bounded by the number of negations in φ, therefore by $|\langle\varphi\rangle|$. Each entry d_i is a multiple $\lambda_{i_1}\lambda_{i_2}\dots\lambda_{i_m}$ of different discounting bases λ_j (there are at most $|\langle\varphi\rangle|$-many such), and since its value must be bigger than ε, the length m of such a multiple is at most $\log_{\lambda_0} \varepsilon$. Therefore the number of candidates for $d_i = \lambda_{i_1}\lambda_{i_2}\dots\lambda_{i_m}$ is bounded by $|\langle\varphi\rangle|^{\log_{\lambda_0} \varepsilon}$; appealing to the fact (see [3]) that $\log_\lambda \varepsilon = \log \varepsilon / \log \lambda$ is polynomial in the length of the description of λ and ε, we obtain the claim. ◀

Proof. (Of Thm. 4.17, sketch) We describe how to avoid the exponential blowup in the translation from $\mathcal{A}_{\varphi,\varepsilon}$ to $\mathcal{A}^{\mathrm{na}}_{\varphi,\varepsilon}$.

Looking at the construction of Prop. 3.4 in case of $\mathcal{A} = \mathcal{A}_{\varphi,\varepsilon}$, we have $V_Q = \{0,1\}$, therefore

$$Q' \;=\; \mathcal{P}(Q \times 2) \times V_\delta \times 2 \;\cong\; \big(\mathcal{P}(Q)\big)^2 \times V_\delta \times 2 \;. \tag{36}$$

Here the original state space Q is bounded by $xcl_\varepsilon(\varphi) \times |\langle\varphi\rangle|^{\log_{\lambda_0} \varepsilon}$, where

$$xcl_\varepsilon(\varphi) \;=\; xcl(\varphi) \setminus \{\varphi_1 \mathsf{U}_\eta \varphi_2 \in xcl(\varphi) \mid \eta(0) < \varepsilon\}$$

is a finite set and the second component $|\langle\varphi\rangle|^{\log_{\lambda_0} \varepsilon}$ is from the proof of Prop. B.2.

The optimization lies in the reduction of $\mathcal{P}(Q)$ that occurs in (36) to

$$\big((Q \times Q^2 \times Q^2) \cup \{\bullet\}\big)^{xcl_\varepsilon(\varphi)} \;, \tag{37}$$

hence from a double exponential to a single exponential; recall from the proof of Prop. B.2 that Q is exponential and $xcl_\varepsilon(\varphi)$ is polynomial, in $|\langle\varphi\rangle|$ and the description length of ε.

The reduction is done concretely as follows. Given a set

$$\big\{\, (\psi, \vec{d_1}),\, (\psi, \vec{d_2}),\, \dots,\, (\psi, \vec{d_m}) \,\big\} \tag{38}$$

of states of Q with a common first component ψ, we suppress the set into the function

$$(\vec{d_1} \wedge \dots \wedge \vec{d_m}) \boxtimes (\_) \,:\; v \;\longmapsto\; \min\{\vec{d_1} \boxtimes v, \dots, \vec{d_m} \boxtimes v\} \tag{39}$$

that does the same job. The latter is a piecewise linear function on $[0,1]$ and hence is presented as a disjunction of pairs $(f_i, [l_i, r_i])$ of a linear function f_i and its domain (here $l_i, r_i \in (0,1)$). Now f_i is represented by some discount sequence so there are at most $|Q|$-many of them. A point $l_i \in [0,1]$ is expressed as the cross point of two linear functions, each represented by a discount sequence. The same goes for r_i. Moreover, disjunction is taken out of a single state in the resulting automaton—from alternating to non-alternating we only need to bundle up states in conjunction. In summary, to express the piecewise linear function in (39) we need: Q to represent f_i; Q^2 to represent l_i; and Q^2 to represent r_i, resulting in $Q \times Q^2 \times Q^2$ in (37).

We consider all those sets in the form of (38), therefore we need $Q \times Q^2 \times Q^2$ for each formula $\psi \in xcl_\varepsilon(\varphi)$. The set $\{\bullet\}$ is in (37) to take care of the case when the set (38) for the formula ψ is empty. ◀

C Reduction of Fuzzy Automata to $[0,1]$-Acceptance Automata

A generalization of $[0,1]$-acceptance automaton is naturally obtained by making transitions also $[0,1]$-weighted. The result is called *fuzzy automaton* and studied e.g. in [20]. Here we show that this generalization does not add expressivity. In fact we prove a more general result, parametrizing $[0,1]$ into a general semiring $\mathbb{K}$ (under certain conditions).

We follow [13] and impose certain conditions on a semiring K of weights.

▶ **Definition C.1** ([13]). A tuple $\mathbb{K} = (K, \leq, +, \cdot, 0, 1)$ is called an *ordered semiring* if $(K, +, \cdot, 0, 1)$ is a semiring, $(K, \leq)$ is a partially ordered set and both $+$ and $\cdot$ are monotonic.

An ordered semiring $\mathbb{K} = (K, \leq, +, \cdot, 0, 1)$ is said to be *lattice-complete* if: $(K, \leq)$ is a complete lattice; the units $0, 1$ of $+, \cdot$ satisfy $0 \leq x \leq 1$ for each $x \in K$; and

$$y + \sup_{i \in I} x_i \;=\; \sup_{i \in I}(y + x_i)$$

for each family $(x_i)_{i \in I}$ and each $y \in K$. We define an infinite sum, as usual, by

$$\sum_{i \in I} x_i \;=\; \sup_{F \in \mathcal{P}_{\text{fin}}(I)} \sum_{i \in F} x_i$$

where $\mathcal{P}_{\text{fin}}(I)$ is the set of finite subsets of I.

A semiring is *locally finite* if the underlying monoid $(K, \cdot, 1)$ is locally finite, that is: for each finite subset $F \subseteq K$, the submonoid of $(K, \cdot, 1)$ generated by F is finite.

The notion of $\mathbb{K}$-weighted (Büchi) automaton is studied in [13], from which the following definition is taken.

▶ **Definition C.2** ($\mathbb{K}$-acceptance (Büchi) automaton, $\mathbb{K}$-weighted (Büchi) automaton)**.** Let $(K, \leq, +, \cdot, 0, 1)$ be a lattice-complete semiring. A *$\mathbb{K}$-acceptance (Büchi) automaton* is a tuple $\mathcal{A} = (\Sigma, Q, I, \delta, F)$, where Σ is a finite alphabet, Q is a finite set of states, $I \subseteq Q$ is a set of initial states, $\delta : Q \times \Sigma \to \mathcal{P}(Q)$ is a transition function and $F : Q \to K$ is a function that assigns an *acceptance value* to each state. We define the language $\mathcal{L}(\mathcal{A}) : \Sigma^\omega \to K$ of $\mathcal{A}$ as

$$\mathcal{L}(\mathcal{A})(w) = \sum_{\rho \in \text{run}(w)} \max\{F(q) \mid q \in \text{Inf}(\rho)\} \;.$$

A *$\mathbb{K}$-weighted (Büchi) automaton* is a tuple $\mathcal{A} = (\Sigma, Q, I, \delta, F)$, where Σ is a finite alphabet, Q is a finite set of states, $I : Q \to K$ is a function assigns an *initial weight* to each state, $\delta : Q \times \Sigma \to K^Q$ is a ($\mathbb{K}$-weighted) transition function and $F : Q \to K$ is a function assigns an *acceptance value* to each state. We define the language $\mathcal{L}(\mathcal{A}) : \Sigma^\omega \to K$ of $\mathcal{A}$ by

$$\mathcal{L}(\mathcal{A})(w) = \sum_{q_0 q_1 \ldots \in Q^\omega} \inf_{n \in \mathbb{N}} \sup_{i \geq n} \big(I(q_0) \cdot \delta(q_0, w_0)(q_1) \cdot \cdots \cdot \delta(q_{i-1}, w_{i-1})(q_i) \cdot F(q_i) \big) \;.$$

These notions specialize to $[0,1]$-acceptance automaton and fuzzy automaton [20] by taking the fuzzy semiring $([0,1], \max, \min, 0, 1)$ as $\mathbb{K}$ in the above definitions.

Locally finiteness of a semiring [13] is central in the following result. Its proof is not hard but the result is not explicit in [13] or elsewhere.

▶ **Lemma C.3.** *Let $\mathbb{K} = (K, \leq, +, \cdot, 0, 1)$ be a lattice-complete semiring and $\mathcal{A} = (\Sigma, Q, I, \delta, F)$ be a $\mathbb{K}$-weighted automaton. If $\mathbb{K}$ is locally finite (Def. C.1), there exists a $\mathbb{K}$-acceptance automaton $\mathcal{A}' = (\Sigma, Q', I', \delta', F')$ such that $\mathcal{L}(\mathcal{A}) = \mathcal{L}(\mathcal{A}')$.*

Proof. Let $(F, \cdot, 1)$ be the submonoid of $(K, \cdot, 1)$ generated by the (finite) set of weights of transitions occurring in $\mathcal{A}$, that is, $\{\delta(q,a)(q') \mid q, q' \in Q, a \in \Sigma\}$. The set F is finite since $\mathbb{K}$ is locally finite. We now define $\mathcal{A}' = (\Sigma, Q', I', \delta', F')$ as follows.

$$Q' = Q \times F \ , \qquad I' = I \times \{1\} \ ,$$
$$\delta'\big((q,k),\, a\big) = \big\{\ \big(q',\ k \cdot \delta(q,a)(q')\big) \ \big|\ q' \in Q\ \big\} \ , \qquad F'(q,k) = k \cdot F(q) \ .$$

The proof of $\mathcal{L}(\mathcal{A}) = \mathcal{L}(\mathcal{A}')$ is straightforward. ◀

It is straightforward that the fuzzy semiring $([0,1], \max, \min, 0, 1)$ is locally finite. This leads to:

▶ **Corollary C.4.** *Let $\mathcal{A}$ be a fuzzy automaton. There exists a $[0,1]$-acceptance automaton $\mathcal{A}'$ such that $\mathcal{L}(\mathcal{A}) = \mathcal{L}(\mathcal{A}')$.* ◀

The main results of [13, 20] concern the characterization of so-called *ω-rational formal power series* over $\mathbb{K}$—those which are generated by ω-regular-like expressions—by $\mathbb{K}$-*weighted* Büchi automata. Lem. C.3 therefore gives us another characterization by $\mathbb{K}$-*acceptance* Büchi automata.